Psychology in society (PINS), 1997, 22, 4-19

A PARALLEL DISTRIBUTED MODEL OF THE PSYCHOLOGICAL DISCIPLINES AND PROFESSIONS

Lance Lachenicht
Department of Psychology
University of Natal
Pietermaritzburg

INTRODUCTION.

Most psychologists in South Africa would admit that a non-reductive model of the relations between individual psychologists and the psychological professions (the *individual interaction* versus *structure* relationship – a central problem of sociology [see Giddens, 1984]) combined with a non-reductive model of the relations between the psychological sub-disciplines and professions would be something well worth having at a time of transformation and reconstruction. A search of the sociological and social psychological literatures will yield few promising candidates for such a model. One exception is the newly developed model of scientific disciplines as (to use Minsky's (1986) phrase literally) *societies of minds*, based upon the model of parallel distributed processing. (Minsky's phrase *societies of minds* is being used literally here because Minsky is actually referring to "agents" internal to the individual that carry out cognitive tasks rather than to individuals who exist in a real society.) Gigerenzer, a cognitive psychologist and an historian of psychology, has suggested that psychologists regularly make use of the "tools to theories heuristic" in which scientists tend to use their tools, which they have come to understand well, as metaphors for the subject they are trying to understand – in the case of psychologists, the mind (Gigerenzer & Murray, 1987).

Gigerenzer has attempted to show the operation of this heuristic in the development of statistical models of cognitive processes, but his clearest example is the way in which psychologists have used the computer as a metaphor in cognitive psychology. Gigerenzer is wrong, however, in supposing that the relation between tools and discipline is simply one-way. The theory and practice of computers has been as much influenced by developments in psychology and neuroscience as psychology has been influenced by developments in computers. The latest extension of the tools to theories heuristic is the use of parallel distributed processing to model social organisations. This model of social processes follows the recent proliferation of computer networks both in academia and in the business world. In a simple computer network each computer on the network is a node, and communication takes place between them in a digitally encoded form (Derfler & Freed, 1993). A new discipline of "distributed artificial intelligence" (DAI) has developed which investigates the principles by which computers that each possess some degree of intelligence can, as a collective, accomplish tasks which no single computer could easily have accomplished on its own (Bond & Gasser, 1988; Gasser, 1991; Hewitt, 1991).

And just as computer science has been influenced by developments in psychology, the theory of parallel distributed processing has been influenced by developments in sociology (Bond & Gasser, 1988; Gasser, 1991; Hewitt, 1991). The relation between the two disciplines is obvious once one realises that parallel distributed processing is fundamentally about dividing up tasks between computers or processors, and that the social division of labour is fundamental to the concept of functional differentiation in sociology. (It is perhaps worth recalling that the concept of functional differentiation is central to the work of Durkheim and Habermas). Introducing an issue of *Byte* dealing with parallel distributed processing, the editor stresses the compelling parallels between the social division of labour and parallel distributed computing, using the division of labour needed to make many pizzas as an illustration:

"Believe it or not, this [the division of labour needed to make a large number of pizzas] is distributed processing. When there is only a little work to do, one person, or processor, can handle it fairly easily. But when the work multiplies, the more hands, or processors, the merrier. In fact, you may find that some people, and some processors, are better at some tasks than others. The benefits add up quickly in terms of cost effectiveness, more efficient use of resources, and quicker response times." (Tazelaar, 1989:212).

Using the idea of parallel distributed processing as a basis for the notion of "distributed artificial intelligence" and therefore as a model of social interaction involving intelligent agents requires extending the cognitive framework so that it can encompass the social. Whilst aware that the cognitive framework has had failures, and that human intelligence far outstrips that of computers, we nevertheless believe firstly that the cognitive framework has had some significant successes (it has enormously enriched our understanding of human thought) and secondly that the cognitive framework has to be viewed dynamically, as continuously solving problems and increasing its approximation to real human intelligence, so that it should not be rejected simply because of some present failures.

There are different meanings that can be attached to the term "parallel distributed processing". A *Byte* article (in the same issue dealing with parallel distributed processing referred to earlier) offers a range of definitions of parallel distributed processing:

"One school of thought considers the client/server model (which uses Structured Query Language [SQL] and transactions from a variety of users) to be a form of distributed processing. Another insists that distributed processing applies only to those systems that attempt to distribute the various tasks or processes of a program across a network to the systems best suited to them. Still another group considers a parallel-processing system to be distributed because it distributes various parts of a program among its own different processors. And then there are the distributed applications and databases." (Wai, 1989:215).

Essential components of the idea of distributed computing, common to all these definitions, are that there be a variety of computational *nodes* (i.e., computers) which are also linked to each other by means of a network of communication channels. In the present article we will distinguish distributed computing from parallel computing. *Parallel computing* typically involves simple nodes of similar kinds communicating with each other in relatively simple ways. For example, in "connectionist" systems of parallel computing each node is a very simple processor (based upon a model of the biological neuron) which updates its level of activation depending upon the nodes to which it is linked and the weights attached to those links (Rumelhardt, McClelland & the PDP research group, 1986). Intelligence in a parallel computer arises from the operation of numerous interacting nodes. But in *distributed computing*, as we understand the term, each computational node has great computational power (it is a computer in its own right), and can communicate in complicated ways with other nodes. In this article we think of distributed computing as involving the distribution of tasks across a network of powerful computers (the second definition given by Wai above). The qualifier "parallel" when preceding "distributed computing" simply indicates that the nodes perform their tasks at the same time.

Distributed computing is essentially focused on how groups of communicating computers can collectively solve problems that are dauntingly complex or even impossible to solve on individual computers. When used as a model of social relations the model must therefore be built around some conception of a problem that a social group is trying to solve. The model will then focus on various ways in which the problem can be divided into sub-tasks and appropriate communication instituted between the computers. Distributed artificial intelligence (DAI) has a wider scope for it sees intelligence as emerging from the interaction between communicating computers (i.e., communication between the computers creates a dynamic cognitive framework, an evolving set of problems, and solutions to those problems). Nevertheless even DAI does presuppose some basic need for cooperation or perhaps some shared enterprise or problem (however general) that provides the foundation for communication between the computers.

PSYCHOLOGY IN SOUTH AFRICA.

Psychology in South Africa is divided in many different ways. It has been divided racially, politically and ideologically. Organisationally, it has been divided into different "divisions" each with an apparently different area of application (though in fact there is considerable overlap). The discipline has been (and still is) divided between different theoretical schools, and it has been divided about appropriate methods. For many outsiders, psychology has a powerful advantage in combining both scientific and practical/applied disciplines. Ideally, the practical disciplines ensure the relevance and value of psychological work; the scientific disciplines help assess existing practice, ensure that new areas of practice are developed, and answer specific questions which arise in practice. Unfortunately, relations between the sub-disciplines of psychology have sometimes been less than cordial, and frequent attempts are made to assert the primacy of one field over the others. (This claim is based upon the personal experience of the author as a teacher and practitioner of psychology and as a member of various psychological groupings; but it would not be hard to demonstrate by examining the record of litigation between psychologists; and by pointing to both the great difficulty experienced in creating a unified body to represent psychologists and the various disputes that have emerged surrounding the registration and categorisation of psychologists).

In fact, the relationships between the different components of psychology can probably be divided into three types (Thagard, 1993): *reduction* (one area is said to be simply a sub-division of another higher and more important area, as in "Industrial Psychology is just a form of applied social psychology"); *residue* arguments ("all versions of problem ***x*** have to be treated / handled / understood by psychologists of type ***y***; anything over and above ***x*** [the residue] is available for other psychologists); and *autonomy* ("our discipline or practice may have the word 'psychology' in its title somewhere but it is really a completely independent discipline or practice and must be conducted without any reference to any other discipline or practice").

Modeling the discipline as a *society of minds*, i.e., as a parallel distributed computer system, shows that this kind of dispute is unnecessary: different sub-disciplines can all be working on the same problem simultaneously, from different perspectives, contributing richness to the overall effort, without any need for reduction, residue or autonomy arguments. The *societies of minds* model allows us to understand the cognitive division of labour within a discipline (and between disciplines), and even suggests that successful disciplines will encourage many different approaches. In what follows we will try to explain how this model works, firstly by sketching how a discipline can be modeled as a parallel distributed computer system in general terms, and then by applying the model to psychology in South Africa in more detail. (Much of the discussion will pertain more directly to the discipline than to the profession, but this is a matter of simplicity of exposition: we believe the same kinds of points could be made about the profession of psychology). From this modeling effort we will turn to looking at some of the general properties of parallel distributed social models, and then to considering some general objections to such models.

ACADEMIC DISCIPLINES AS PARALLEL DISTRIBUTED COMPUTATION.

Several researchers in DAI have used academic disciplines as examples of parallel distributed networks so there is no novelty in applying this model to psychology in South Africa, though there may nevertheless be value in spelling out some of the properties of the model for this particular case, and for those who are not familiar with this style of modeling. Essentially, the scientists in the discipline (in this case psychologists) are viewed as *nodes* in a *communication network*. If we are to view psychologists as nodes in a computational network then we have to view them as very powerful computers capable of communicating in natural language, preparing written documents, making drawings, interacting with clients, and also of deciding when and how to interact with the network of other psychologists of which they are a part. As with all computational networks, the network of psychologists is concerned with transmitting information. The routes by which information is transmitted in a profession and an academic discipline such as psychology include such direct methods as personal contact, teaching and referrals, and such relatively indirect but much more widespread contact as journal articles, book publishing and contact with the general media. All of these methods of transmitting information are complex and require considerable amounts of encoding and decoding in order to work well. The encoding and decoding operations are not always transparent, so that information transmission is not perfect. And, as with all networks, it takes longer and is relatively more costly to transmit information over long distances than over short distances.

Trying to flesh out this model we can see that the psychological professions and disciplines can be broken into relatively small clusters which are tightly inter linked by fast communication channels, and these clusters are more loosely linked to larger groupings. Clusters and groupings in which most psychologists can be found include: (i) *Student and teacher clusters*, which, particularly for postgraduate teaching, are often very tightly linked. The links may also endure over a very considerable period for the postgraduate students belonging to the same class. (ii) *Collaborator clusters*, of people working in the same project or problem in the same institution, who communicate frequently. (iii) *Colleague groupings* consist of people from the same or related sub-disciplines, who work in the same institution, who are not collaborators, but who nevertheless exchange information fairly frequently. (iv) *Acquaintances* are groupings of colleagues who meet at irregular intervals at conferences, and who perhaps refer to each other and review each others' work. Finally, (v), *correspondents* are loose groupings of psychologists whose principal form of contact is through letters and e-mail.

In all the groupings we have just considered the networks are based upon a degree of direct contact, or the direct exchange of information. But disciplines and professions also make use of the indirect exchange of information. The most obvious indirect communication channel is journal articles. Others include book chapters, books, presentations at conferences, funding proposals, referrals, reports, journal refereeing, and the like. Each of these channels has different speeds of operation and different degrees of penetration into the network. A journal article may take years to appear in print. Presentations at conferences are more rapid. However, journal articles will probably reach a wider audience than conference presentations, though this depends upon the choice of the journal. Part of the intelligence each "computer" in a social network has to have is the ability to select the appropriate means of dispersing information across the network. Psychologists who are carefully planning their careers will attempt to select journals which have greater prestige and wider readership. But it is also part of the beauty of this kind of model that it can account for the psychologist who chooses to sit on his or her findings for years before publishing them.

We have concentrated upon the links within psychological disciplines and professions, but of course the social network also has links between sub-disciplines and professions, and between completely different professions and disciplines. Most psychologists read material from more than one area of psychology, and many read material from related disciplines such as medicine and health or zoology, or philosophy or labour relations. In professional contexts psychologists may team up with other professionals such as psychiatrists, surgeons, dentists, lawyers, engineers, managers, trade union representatives, and statisticians. In some cases, these links may be tighter, and transmit more information, than the links between mere colleagues of the same profession from the same institution.

As many philosophers of science have noted, all sciences are based upon trust. No single scientist can repeat every experiment that has been carried out by others, and scientists simply have to take the greater part of their disciplines upon trust. This trust seems to be based on shared framework in which the discipline is seen as something intrinsically valuable and this shared value provides the basis for being part of the network. Similarly, professions are based upon a basic presumption of integrity and honesty which in turn seems to be based upon a shared framework that views the service the profession offers as valuable and worthwhile. Nevertheless, an issue of great relevance to any academic discipline or profession is conflict between members of that discipline or profession. Surprisingly, it seems that DAI models are capable of offering some explanation of conflict within a network. All real computer networks depend upon procedures for handling conflicts. For example, if a problem is divided into portions each of which is handled by a different computer on the network, all the computers will initially start with the same information. But as some computers achieve results the information which others are using may become outdated and there have to be procedures for eliminating such discrepancies. (Some of these difficulties are discussed in Pountain & Bryan, 1992). Very real conflict also occurs over access to the limited communication channels, and networks depend upon procedures for mediating such conflicts and sharing the resources between the computers (Derfler & Freed, 1993).

In the case of human social networks, such as those we find in academic disciplines and in professions conflict may be very real. Conflict may occur, for example, over acceptance of articles into journals, with many articles being rejected or requiring modification. Negotiation procedures often arise in such contexts, with the author attempting to justify his or her position in the light of the referees' remarks. Applications for funding are another source of conflict, in which the applicant may need a good understanding of the kinds of research which are likely to receive funds. In more applied contexts, conflict may arise when professionals disagree about a particular case. Professionals may also disagree about who has the right to offer particular services (an issue of access) and under what circumstances clients must be referred to other professionals (an issue of priority). And just as issues of priority and access to limited resources trouble real computer networks so issues of priority and access are a constant source of conflict with academic disciplines and professions. In human social networks the intellectual requirements for mediating conflicts may be very demanding, requiring sophisticated reasoning and modeling skills, for opponents will need to build models of each other and of the requirements for success (Thagard, 1992).

A NON-REDUCTIVE MODEL AND GROUP RATIONALITY.

It might be argued that a model based upon parallel distributed computing must necessarily be reductive, either of persons or of social structures. Actually, distributed computing involves networks (the discipline) of highly capable computers (analogous to persons) which do not simply take orders from other computers on the network. However, the network has a real and independent existence apart from any of the computers on the network, and any particular computer can be replaced without damaging the network. Thus both the computers and the network have independent existences. The network controls flows of information between the different computers, and the different computers have to react to (and perhaps translate) that information and supply their own information which the network will then distribute.

Leydesdorff (1993) attempts to show the operational independence of the social system from individual action in the parallel distributed network model of social organisation by pointing out that the communication system (the network) adds uncertainty to communications which cannot be attributed to any of the actors (i.e., computers) involved. In addition, he contrasts such a proof of the independence of the social system with the "double contingency" of ego and alter arguments normally employed by sociologists, and suggests that the network argument has the advantage over the traditional arguments of clear independence of the intentions of the communicators involved. Leydesdorff's argument proceeds as follows:

"In the mathematical theory of communication ... the expected information content of the system is by definition equal to the uncertainty within the system. As long as an actor processes on its own, it (as a system) can only attribute its uncertainty to itself. However, as soon as two actors communicate, neither of them can internally generate the information necessary to exclude the possibility that the uncertainty has originated from (noise in) the communication. The uncertainty can be attributed to each of them *or* to the communication. Therefore, one has to assume that the communication system itself is able to generate uncertainty by operating." (Leydesdorf, 1993:58-59).

A consequence of the operational independence of the network is the possibility that a distributed network may have properties which differ from the summation of the properties of the linked computers. In other words, proofs of non reduction hold the promise of demonstrating that group rationality is greater than any particular individual's rationality. This suggestion has been made by Paul Thagard (1988). He suggested that viewing scientific communities as heterogeneous processors (i.e., computers using different information, programs) operating in parallel could enhance scientific progress. (The argument is very similar to the evolutionary and population dynamics arguments of Kitcher, 1990). The basic idea of these writers is that by having many competing scientists with different starting points and different programs, the scientific enterprise as a whole can be viewed as a selection process (selecting the best theory or findings), similar to evolutionary selection. It is not difficult to generalise such arguments from scientific disciplines such as psychology to professions. For example, most professions consist of competing individuals with the profession as a whole being served by the selection of those services (and those professionals) that best meet some public demand.

Viewing science or the professions as a parallel distributed computational process, however, makes it clear that one does *not* have to start with premises asserting the radical individuality of the different scientists or professionals, for the very operation of the network will ensure that not all scientists are operating from the same information. Networks may be sparsely connected in some areas, with only very devious routes between any two particular nodes. Thus ideas or data generated at one node will not immediately be communicated to all other nodes. In the sciences and professions network sparseness may be the result of national and institutional factors, as well as a result of human limits on information processing (e.g., the fact that no one can read everything or talk to everybody). Secondly, in human networks there is no clock governing the transmission of information. Some scientists rapidly transmit their findings, some sit on them for a great length of time. Transmission in human networks can also be very slow. It can take years for research to appear in a journal, and it can take months for even habitual readers of the journal to actually read the information. Still another point is that information transmission in human networks tends to be fragmentary: not all the information a scientist knows is encoded into a written article, and not all the information in the article is decoded by its readers. Networks clearly do not transmit information perfectly, for a variety of reasons.

The heterogeneity of the nodes in a network is not necessarily the result of ineradicable individual differences, for the impediments to information transmission across the network will ensure that the nodes start with different information. Each node will also communicate habitually with a different set of nodes. Not surprisingly scientists or professionals who start with different information (training), who communicate with different groups, and who receive different information will make different decisions and arrive at different conclusions. All of this is sufficient to provide the basis for a "variation and selection" mechanism by which the network as a whole may choose the best available theories and data.

And indeed, simulations may show that parallel distributed networks can be much more efficient problem solvers than isolated computers. A recent simulation (Clearwater, Huberman & Hogg, 1991) found that a group of cooperating agents engaged in problem solving can be much more effective than either a single agent or the same group of agents working in isolation from each other. But of course, such simulations may only be showing a part of the picture, for there is a large literature to suggest that sometimes group problem solving may be less than the sum of the individual problem solving capabilities!

Apart from providing a non-reductive model of scientific and practical disciplines, the model of science as *societies of minds* offers insights into the functioning of these disciplines. Both socially and cognitively science involves a tension between cooperation and competition. Competition and cooperation involve both individual attitudes, biases and knowledge but can also be mediated and influenced by the flow of information through a network (where that information is unlikely to be perfectly distributed, especially if the network involves some sparsely connected areas: it is more likely to form local nodes). The metaphor of *societies of minds* connected by a network can enable us to understand both how knowledge is social without neglecting the contribution of individuals and individual cognition in its development.

HIERARCHY AND THE FLOW OF INFORMATION.

An additional and important advantage of the parallel distributed model is that it enables the integration of computational models of individual cognition (which, although not perfect, have enormously enriched our understanding of human thought) with an analysis of communities of practitioners and scientists understood in terms of distributed computation, so that we can start to see how sociological and psychological theories can be integrated. Sociologists such as Leydesdorff have shown that the model is particularly rich from a sociological perspective, so again, we need have no fear of mindless reduction of one discipline to another.

Among the sociological insights derived from the model of parallel distributed computation which Leydesdorff considers are demonstrations of how the network can limit the options of any particular node (i.e., how society can constrain an individual's choice through the creation of "situations") and of how hierarchies can be modeled in networks. Here we will briefly review the creation of control structures and hierarchies in networks.

Consider a network with two local nodes ***A*** and ***B***. ***A*** and ***B*** are directly connected, but the rest of the network, however complex, also serves as an alternative parallel link between ***A*** and ***B***. The situation is depicted in the Figure 1.

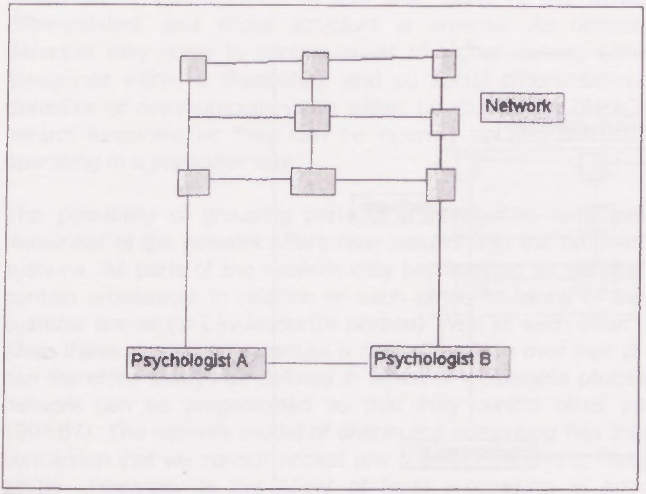

Figure 1

Now there is an electrical engineering theorem (the theorem of Thevenin) which would enable us to calculate the impedance of the entire alternative route from ***A*** to ***B*** (however complex) and therefore think of it as a single complex node, **C**. In a social network, this would be equivalent to replacing the entire network with the sum total of the uncertainty of all its parts. This can be depicted as in Figure 2.

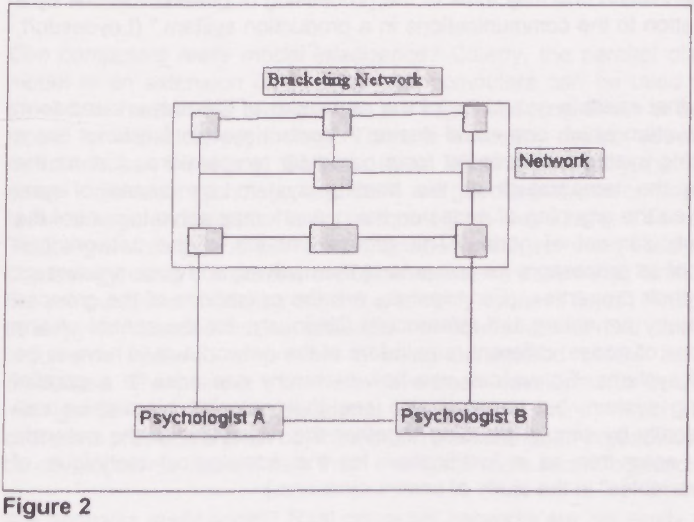

Figure 2

But once we have started grouping the network into complex nodes in this fashion it is possible to decompose it into smaller sets of complex nodes as well. So it should be possible to decompose the complex node ***C*** into two smaller nodes ***C'*** and ***B'***. This is depicted in Figure 3.

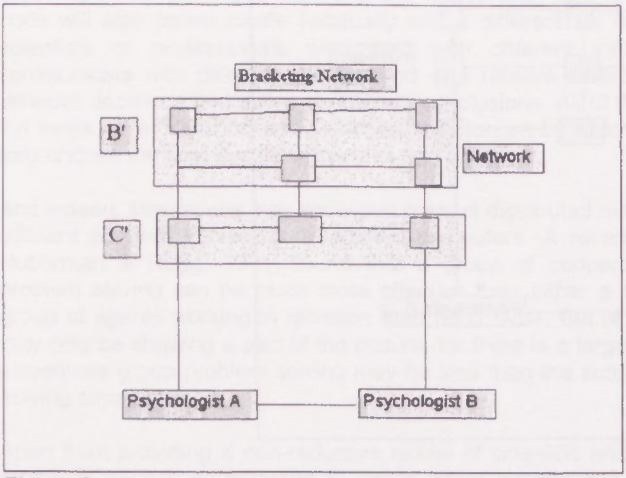

Figure 3

Now, depending upon the grouping rules, it is possible to view ***C'*** and ***B'*** as "localised processors that are performing their own transformations on the signals carried over their respective lines. For example, one may think of ***B'*** as a processor which stabilises the link by damping deviations from a preset value, while ***C'*** modulates variance. Or in terms of a social organisation one may think of the processing of communications by a control system in relation to the communications in a production system." (Leydesdorf, 1993:64).

It is not hard to see that suitable groupings of the remainder of the network can form any type of control system which one could desire. A control system functions like a thermostat in a heating system which is set for a particular temperature. Just as the thermostat regulates the temperature of the heating system, by means of goal referential feedback, so the grouping of nodes on the network may serve to control the output from any particular set of nodes. The grouped nodes of the network can therefore be thought of as processors (or computers) themselves, and such processors can be analysed for their properties. (For example, are the operations of the grouped nodes recursive, thereby permitting self-reference?) Obviously, for the control of any particular combinations of nodes, different remainders of the network would have to be grouped into control systems. So we can see how hierarchy can arise in a parallel distributed processing system, but we can also see that complex hierarchies can always be studied locally by simply grouping together the remainder of the network. (Leydesdorff *op cit.* sees this as a justification for the sociological technique of "bracketing context variables" in the study of communications.)

Social hierarchy in a network model can thus be seen as localisable *densities of communication* within a network. Such densities may undergo transformations which affect the operation of the network within them, and thereby their function for the remainder of the network. In this way, parts of the network become functionally differentiated, and social structure is created. As differentiation proceeds, local densities may come to contain areas of higher density within themselves (like sub-disciplines within a discipline), and so social differentiation proceeds. These local densities of communication can either be studied as black boxes serving particular control functions, or they can be opened up and studied as a detailed network operating in a particular way.

The possibility of grouping parts of a network to form control processors for the remainder of the network offers new insights into the problem of hierarchies in social systems. All parts of the network may be rewritten as parallel connections which may contain processors in relation to each other. In terms of the network, these control systems are all (in Leydesdorff's phrase) "next to each other". Hierarchy occurs only when these processors exercise a control function over part of the network. "Hierarchy can therefore always be defined in terms of localisable processors: some parts of the network can be programmed so that they control other parts of it." (Leydesdorf, 1993:67). The network model of distributed computing has therefore brought us to the conclusion that we cannot accept any *a priori* concerning hierarchies. In Leydesdorff's words, "hierarchy is the effect of local processing in an otherwise heterarchical network." (Leydesdorf, 1993:67).

SOME DIFFICULTIES.

Theoretical problems still remain with the distributed artificial intelligence model, and we will consider some of them here, but we wish to note (1) that given the dynamic nature and rapid development of cognitive science and artificial intelligence they may be overcome in the future and (2) that there are few non-reductive alternative models. Most of the difficulties we discuss are drawn from P. Thagard's (1993) discussion of the limits of DAI modeling.

Can computers really model intelligence? Clearly, the parallel distributed processing model is an extension of the idea that computers can be used to model the mind. Indeed, the entire idea presupposes the cognitivist programme of using computers as a model for individual psychology. There is a healthy debate about the possibility that computers can think (e.g., Collins, 1990), and the possibility that they can be used to model individual psychology, and there are good arguments on both sides. This debate is far to large and complex to be adequately reviewed here. It seems sensible to adopt Thalberg's strategy and admit the difficulties which surround the modeling of individual psychology using computers, but to nevertheless assert that "artificial intelligence and the computational modeling of mind are highly dynamic fields" (Thagard, 1993:59). This strategy is intended not to establish the soundness of DAI but to counter *a priori* arguments intended to prevent it being considered before its case can be fully articulated. The point is that there have been developments in these fields (artificial intelligence and the computational modeling of mind) which have helped them overcome previous *a priori* objections. On the balance of probabilities there are likely to be developments that will help them overcome more recent objections.

Are networks really social? Real computer networks are not really social in the way in which humans are social. The interactions between computers linked on computer networks are enormously primitive compared to those amongst people. Nevertheless, some analogues of social problems do appear in computer networks. Computers do use different representational schemes, different inference engines, and it may be difficult to translate one kind of information from a particular computer into a form which another computer on the network can use. All of this means that simply pointing to the difficulties and complexities of human communication is not enough to rule out the use of network models of social organisation. In his book Collins (1990) specifies three ways in which science is essentially social. Thalberg considers each of these and attempts to show that a DAI system could be used to model that aspect of the social nature of science. Computers on a network could, for example, be used to check the validity of hypotheses and beliefs just as people in a scientific system do, for each computer could bring to bear its own representations and programs in checking the validity of any hypothesis passed to it from the elsewhere on the network.

Similarly, Collins' notions that scientific conclusions are matters of social consensus can be modeled by the notion that processors on the network will pass information back and forth until some state of equilibrium about a particular hypothesis is reached. Collins' final point is that the transfer of scientific skills requires social interaction, for most scientists learn their skills in an apprenticeship with a more skilled practitioner. (This is certainly true of psychology of all persuasions). Thalberg notes that:

"Collins' point does not, however, undermine the DAI perspective, for it can be interpreted as showing the need for some kinds of communication to be particularly intense. Intensity can be a matter of amount of information transmitted – articles typically report much less than the experimenter knows – and format of information ... There is no reason in principle, however, that AI processors equipped with capacities for visual representation and interaction with the world could not communicate skills in the complex ways that Collins describes. Whether AI will accomplish such communication will depend upon the success of the whole research program and cannot be decided *a priori*." (Thagard, 1993:61).

However, given the above attempt to counter *a priori* refutations of the DAI programme on social grounds we have to acknowledge that one aspect of human social life clearly cannot be captured by computer networks as presently understood: the use and the threat of physical force, i.e., such social phenomena as imprisonment or censorship or murder. And clearly the use of, or the threat of the use of, force may underlie the creation of social hierarchies that themselves cannot be justified on the basis of the network of individuals and the flow of information. However, the fact that DAI models are able to offer a model of group rationality does mean that they can be used to highlight and offer a critique of the arbitrary use of force.

Methodological individualism as a social research programme. We have been at some pains to stress that DAI models are not reductive, and offer separate and real existences for both individuals and society. However, such non reductionism clashes with the philosophy of methodological individualism which is widely employed in the social sciences. Essentially, the philosophy of methodological individualism asserts that all attempts to explain social and individual phenomena must refer exclusively to facts about individuals. The present paper is not the correct context in which to explore the strengths and weaknesses of this philosophy. Suffice it to say that distributed processing models may offer a viable alternative by showing that non reductive approaches are possible.

PRACTICAL IMPLICATIONS.

Perhaps of greatest importance to the reconstruction of psychology in South Africa are the practical implications of the parallel distributed processing model of social organisation for psychology nationally. There are a number of objections to drawing practical implications from DAI models as set out in the present paper. Firstly, it may seem premature to be drawing practical implications from a model so recently developed, and still so contentious. But there are not that many alternative models from which to draw practical implications, and the implications seem to be fairly sensible and useful. Secondly, a more important limit on practical implications drawn from the present paper arises from the very general level of presentation: we have been trying to show the possibility of creating DAI models of social groups and have concentrated on establishing general principles and offering defenses against *a priori* objections. Much more detailed DAI models will be needed for practical purposes. These may well take the form of simulations. An example of such distributed simulations (in an economic context) may be found in van der Wal (1996).

A general message for psychology arising from a consideration of distributed processing models is that we should pay particular attention to the links between individuals and institutions dealing with psychology in South Africa. Much attention has been focused upon understanding the different cognitive starting points of disparate psychologists. But it is just as important to pay particular attention to the flows of information between individuals and between sub-disciplines We may greatly strengthen our discipline simply by fostering much more frequent and better quality contact. Here attention should be paid both to direct and to indirect means of contact, including conferences and journals, but not excluding the contact which occurs in classrooms and between institutions. Enhanced international contact will also be very beneficial for the discipline in South Africa, and local psychologists should be encouraged to travel overseas.

A second implication is that it is important to welcome diversity in opinion, approach, theory and practice. Social networks, in the end, serve as selection mechanisms for the best opinions, approaches, theories and practices. Having sufficient diversity of approach in the network allows the network to act as a selection agent, filtering out and transmitting the best procedures across the network as a whole. Of course, the need for diversity has to be counterbalanced by the need for concentration of network resources on some particular set of problems. Relatively small networks, such as exist in South Africa, faced with too great a range of problems, can easily be reduced to the equivalent of a set of isolated nodes each working on a unique problem. In part, greater international contact will reduce this danger, for the isolated nodes can, as a result of such international contact, at least form some connections to related overseas nodes and networks; but it does seem sensible to find some problems where the collective resources of South African psychologists can be focused so that the synergies of belonging to a network of intelligent processors working on related problems can be harvested.

The distributed processing, or *societies of minds,* model suggests that we should reject attempts to discuss the relationship between the various components of psychology in South Africa in terms of *reduction*, *residue* or *autonomy* arguments. It will be far more useful to consider the actual network contacts and the actual flows of information between the various components of the discipline and associated professions. One can ask questions such as: Are the links which bind counseling and clinical psychologists to their professions really very different? Perhaps, in asking such questions, it is possible to find that some parts of the discipline or psychological professions really do have so little contact with other components that they can rightly be considered autonomous. But if this is the case, it may be just as sensible to attempt to increase the contact between these components and integrate them more tightly into the overall discipline as to acquiesce in separation. Reduction arguments about the relation between components of psychology do not make sense in the light of a network model. We have attempted to show that all hierarchy in networks is local in character. (We acknowledge that such hierarchies may in fact be based upon non-rational force). It is simply not sensible to argue the priority of certain positions over those of others. Rather, an effort should be made to interconnect networks of people working on similar problems so that a synergy between the different efforts can be achieved. Residue arguments which involve claiming some area of psychology as uniquely belonging to a particular component of the discipline depend upon some notion of power, notably the ability to restrict others from entering a particular domain. We have seen that networks are fundamentally heterarchies; if power differences occur in a network, then this is because local densities have arisen in the network, but the power such densities confer is essentially local in character. Rather than attempting to fight about the rights and wrongs of such power differences, it seems better to attempt to understand the way in which such local densities in the network have arisen, and to understand the character of the control they are attempting to exert over the remainder of the network. Fruitless arguments about the proper function of any particular component of the network should be avoided. But remember that any network model consists not only of controlling local densities but also of intelligent processing nodes which can choose to resist any control which cannot be justified.

In South Africa the legal framework governing psychological professions creates some structures that, since they are based upon force rather than network considerations, clearly cannot be modeled by distributed models. However network structures probably include the various psychological societies, teaching and research institutions and journals. We need analyse very carefully why so many psychologists do not participate in these "local densities" and resist them. (For instance less than half the registered psychologists in South Africa have joined PsySSA, the new unified psychological society). Since networks are fundamentally heterarchies, we should remember that any hierarchies which do exist are nevertheless fairly fluid, and can be reformed and restructured simply by building new channels of communication between relatively isolated groups of nodes, and between local densities of nodes and other nodes. Hierarchy did not exists prior to the network: it is a local product of the network, and changing the network will change any existing hierarchies. In a deep sense, that is exactly what the present reconstruction of psychology is about.

A final point that arises from the distributed computing model concerns the importance of a shared "problem" or framework that justifies network participation and prompts individuals to become part of the network. Practitioners of a discipline or a profession must share a concern or a problem with others like themselves to become part of the network. This suggests a number of non-structural reasons as to why network participation may be limited or absent. (Structurally, alternative network links may provide individuals with adequate support and information). Perhaps the registration system in South Africa that does not require continuing education in the profession means that many professionals do not have a need to become part of the network. Perhaps the profession or discipline really is so ill-defined that there really is no common problem that all practitioners share thus rendering participation irrelevant. Or perhaps the practitioners have grown cynical and have lost their trust in other scientists or their belief in the integrity of other professionals so that participation is worthless. But whether there is any truth in such suggestions is perhaps less important than realizing that the distributed computing model reminds us that disciplines and professions, however ruthlessly competitive they may be, ultimately depend upon shared problems and values. Without them we have nothing to say to each other. It is of the utmost importance for our discipline and our professions that we foster a shared trust and a shared perception of integrity.

REFERENCES.

Bond, A & Gasser, L (1988) (eds) **Readings in distributed artificial intelligence**. San Mateo: Kaufamn.

Clearwater, S, Huberman, B, & Hogg, T (1991) Problem solving by committee. **Science, 254**, 1181-1183.

Collins, H (1990) **Artificial experts: Social knowledge and intelligent machines**. Cambridge, Mass.: MIT Press.

Derfler, Jr. F J, & Freed, L (1993) **How networks work**. Emeryville, Calif.: Ziff-Davis Press.

Gasser, L (1991) Social conceptions of knowledge and action: DAI and open systems semantics. **Artificial Intelligence**, **47**, 107-138.

Giddens, A (1984) **The constitution of society**. Cambridge: Polity Press.

Gigerenzer, G & Murray, D J (1987) **Cognition as intuitive statistics**. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hewitt, C (1991) Open information systems semantics for distributed artificial intelligence. **Artificial Intelligence**, **47**, 79-106.

Kitcher, P (1990) The division of cognitive labor. **Journal of Philosophy**, **87**, 5-22.

Leydesdorff, L (1993) "Structure" / "Action" contingencies and the model of parallel distributed processing. **Journal for the Theory of Social Behaviour**, **23**, 47 - 78.

Minsky, M (1986) **The society of mind**. London: Heineman.

Pountain, D & Bryan, J (1992) All systems go. **Byte, 17**, 112-136.

Rumelhardt, D, McClelland, J & the PDP Research Group (1986) **Parallel distributed processing: Explorations in the microstructure of cognition**. Cambridge, Mass.: MIT Press / Bradford Books.

Tazelaar, J M (1989) Distributed processing. **Byte**, **14**: 212.

Thagard, P (1988) **Computational philosophy of science**. Cambridge, Mass.: MIT / Bradford Books.

Thagard, P (1992) Adversarial problem solving: Modeling an opponent using explanatory coherence. **Cognitive Science**, **16**, 123-149.

Thagard, P (1993) Societies of minds: Science as distributed computing. **Studies in the History and Philosophy of Science**, **24**, 49-67.

van der Wal, R (1996) Designing a distributed simulation game. In **Games, Graphics, Multimedia. Dr Dobbs Journal, 21(15),** 21-26.

Wai, G (1989) Take your pick. **Byte**, **14**, 215-223.