

THE POWER OF PARALLEL THINKING

Michael D. COHEN

University of Michigan, Ann Arbor, MI 48109, USA

Received August 1981, final version received March 1982

A small computer model demonstrates that an appropriate organization of boundedly rational individuals can find optimal policies in an environment that is overwhelmingly complex for unorganized decision makers. The model is also used to identify conditions under which optimal — or even good — policies are not found. The demonstrated adaptive power of the model is interpreted in light of recent developments in the theory of computational complexity that place new stress on powerful methods of search, and of new models from computer science which markedly advance search effectiveness by harnessing parallel structures of information processing.

1. Introduction

The ability of organizations to adapt to the demands of their environments is an issue of major significance for the social sciences. As Simon (1969) has pointed out, the limits of that ability will be closely related to the approximation adequacy of rational models of organizational action. Given its importance, and the extensive attention that has been directed to organization-environment relations [Starbuck (1976), Pfeffer and Salancik (1978), Aldrich (1979)], it is surprising that we do not have a satisfactory account of how an organization could adapt successfully to an environment of any complexity.

One can, of course, simply posit that adaptation never occurs or is always perfect, but neither of these implausible courses seems promising if the limits of organizational adaptation are to be studied rather than settled by assumption. To explore the middle ground we need a theoretical account of adaptive organizational decision making that is both plausible in the face of what we know about organizations and their component actors, and demonstrably capable of adapting to a complex environment. It cannot violate the former condition by endowing the organization or its members with unrealistic volumes of information or powers of inference, for that would cast doubt on its fidelity to our hard-won existing knowledge of real organizations. It must meet the second condition or there is little reason to credit its patterns of limitation or pathology. If organizations do sometimes

adapt to complex environments, why believe results from a model that never could? This is essentially the test of 'sufficiency' which Newell and Simon (1972) argue deserves a central place in theory development.

This paper presents a computer model of an organization that satisfies these two partially antagonistic conditions. It incorporates the major features and limitations of organizational actors and structures usually found in contemporary organization theory, and it adapts with remarkable success to environments of very considerable complexity. It is able to do this because the model's processes and structures give it an adaptive capacity that dramatically exceeds the simple summation of the capabilities of its component individuals. In model variants in which these interindividual processes are disabled or inappropriately arranged, adaptive abilities are orders of magnitude weaker.

The models to be reported here are members of a very large family that can be studied using a system that has been developed for constructing computer models of organizations. Within this system one can study the effects of variation in task decompositions, patterns of incentives, level and types of environmental uncertainty, information flows and other organizational factors. The method captures much of the rigor of mathematical theory and much of the richness of verbal theory. These gains are purchased at the price of increased difficulty in the interpretation of results due to greater model complexity. For many purposes, however, the trade appears to be an attractive one.

The modeling system provides modular individuals with sharply bounded information processing capabilities. Thus it is built upon the fundamental revelations of empirical studies of decision making [Estes (1978), Anderson (1980)]. The model individuals can be assembled into virtually any organizational structure that one wants to examine and provided with virtually any pattern of incentives and rewards for their decisions. In effect, such a specification of rewards flowing from alternative choices creates an environment to which the organization will adapt or, we might say, a set of interrelated problems for which the organization will search for solutions. Other factors can also be set as an investigator's interests dictate. These include: friendship networks, agenda for meetings, rules for making collective decisions, noise, lags, or other environmental uncertainties, and organizational precedents. Only a very few of these capabilities are exercised in the work reported here, however.

Once the model is started it takes only a few seconds of computer time for a complete case history of organizational decision making to unfold. Model meetings are held. Decisions are made. Results of new decisions are evaluated by model individuals whose satisfaction state may change as a result of changes in current policy. The model produces interactions among colleagues called 'talk' processes and these may lead to the proposal of further policy

changes. Incentives of model actors may change. New meetings are held, and the cycle is repeated for a few dozen periods. The case can be recreated exactly by rerunning the model with the same random number seed. The sensitivity of results can be investigated by intervening in mid-run to change particular events or by making multiple runs with different random number seeds. Theoretically interesting patterns can be extracted from data on sets of cases generated by systematically varying organizational parameters.

This very quick sketch requires elaboration. The next section gives a fuller account of the system developed for modelling organizational processes. After it comes a discussion of the nature of environmental complexity and a description of the model environment studied here, a transportation problem originally posed by Elantzig. This is followed by an account of the parameter settings that were used to create three specific organizational models, a basic model and two variants. As each of the models tries to cope with ten variations on the model environment, they generate thirty case histories of organizational decision making. The succeeding section reports the cases and the patterns they present of success and failure in finding optimal organizational policies. The final interpretive section connects the results to recently developing literatures on heuristic approaches to computational complexity and on the behavior of systems of parallel processes.

2. Major elements of the modelling system

The models employed below are, as noted, elements of a very large set that can be generated by varying elements of the organizational modelling system. The general system is a computer program and subroutines that encode organizational structures and processes. By altering data that initializes the program or by small changes in the subroutines, wide variations can be induced in organizational structures and processes.¹ The major elements of the system fall into four clusters.

2.1. The policy string

At any moment the policies in effect in a model organization are represented by the state of a sixty-four bit string. The most natural interpretation is that each bit displays whether (on) or not (off) a particular standard operating procedure of the organization is currently in force. For some applications very different interpretations may be convenient. However, the favored interpretation meshes comfortably with an evolutionary view of

¹The system is in FORTRAN, and is transferable to most computer installations with relatively minor adjustments. Complete detail is available from the Institute of Public Policy Studies, University of Michigan, as discussion paper no. 151, "Documentation of an Organizational Modelling System".

organizational policy in which, following Winter (1975), routines of the organization are treated as analogous to biological genes. An individual bit position will often be referenced as a 'locus'. The whole string will be called 'the policy'.

2.2. The environment function

The meaning for any individual in the organization of a particular overall policy is determined by the interaction of the policy and the organization's environment. In each time period each individual's identification number and the current policy become arguments of an environment function that returns to the individual the actual rewards he or she experiences in that period. This is called operating 'in real mode'. The environment of the organization may be something different for every individual and in every time period. The only theoretical limit is set by what can be written as a mathematical function from the space of possible policies, individual identification numbers, and time periods, into the rational numbers. In the case resembling the simplest firm models traditionally used in economics, all individuals may have the same environment function, typically one returning the utility associated with a profit stream produced by the policy. However, the system permits the study of more complex cases in which utility functions differ among individuals, and/or change over time, and in which environmental responses to organizational policies depend on past policies.

The environment function may also be called by an individual in an *estimating mode*. In a typical model constructed in the system there will be many such calls per time period. In this mode the function returns to the individual an estimate of the rewards that would flow from some altered version of current policy. In the limiting case in which individuals can perfectly foretell the consequences of alternatives, the estimating mode may be identical to the real mode. The more interesting cases, however, will be those in which the estimates of rewards are distorted or erroneous assessments of the actual consequences of a potential policy.

2.3. Attributes of model individuals

An individual in models built with the system is represented by a small set of ideas for policy improvements, an aspiration level, a satisfaction state, and a just-noticeable-difference (JND). In effect, ideas are represented as sixty-four bit strings from the alphabet $\{0, 1, \#\}$, where $\#$ means 'don't care'. An idea thus specifies a subspace of the space of possible policies. A typical idea for an individual will be blank ($\#$) on most, but not necessarily all, loci outside his or her domain of formal responsibility. This is due to the very limited attention range and sharply bounded cognitive power implemented

in the model individuals. Operating in isolation, individuals are capable of systematic exploration of alternatives only in their domains of formal responsibility and only if those domains are relatively small — no more than about half-a-dozen loci. Ideas about loci outside the domain of responsibility mostly come to an individual from others via organizational processes described in the next section.

When an individual considers an idea, either one self-generated or one suggested by someone else, it is evaluated by substituting current policy for all the don't care (\neq) loci and then calling the environment function in estimating mode. This amounts to evaluating the idea by reference to the element in the idea subspace nearest to the location of current policy. In the process of retaining or discarding new ideas only ordinal use is made of the resulting estimate. If the new idea has a better estimate than the worst idea in an individual's retained set of 'good' ideas, the new idea replaces the worst member of the set. Otherwise the new idea is forgotten. The retained set of good ideas is quite small, ordinarily about five.

Each individual also remembers an idea that incorporates the best setting so far encountered of the loci for which the individual is responsible, assuming current policy to prevail on other loci. The process of considering a possible new idea is illustrated in fig. 1.

The JND is the difference between the evaluation of the best and worst members of the set of good ideas, divided by the number of good ideas less one. The aspiration level is a weighted average of recent period rewards. The satisfaction state rises when rewards exceed aspirations by at least a JND and falls when the opposite event occurs. Both high and low extremes of satisfaction dampen somewhat individual exploration of alternatives to current policy to capture the effects of satiation and discouragement.

2.4. *Attributes of model organizations*

The organizational structure of any model developed using this system is determined by five major components: the pattern of individual formal responsibilities; the regular meetings held in the organization together with their agenda; the method of calling irregular or special meetings and setting their agenda; the processes of give and take over ideas that occur as individuals interact in meetings; and the procedures in the organization for making authoritative choices of policy. The first and last of these can be altered by varying the initializing data set. Meetings and agenda setting can be changed by altering subroutines. Idea exchanges in meetings, which will be called 'talk' processes, are relatively invariant program features.

The responsibilities of each individual are the loci to which he or she predominantly attends. This number is limited — usually it is about four. Within that small domain individuals examine deviations from current policy

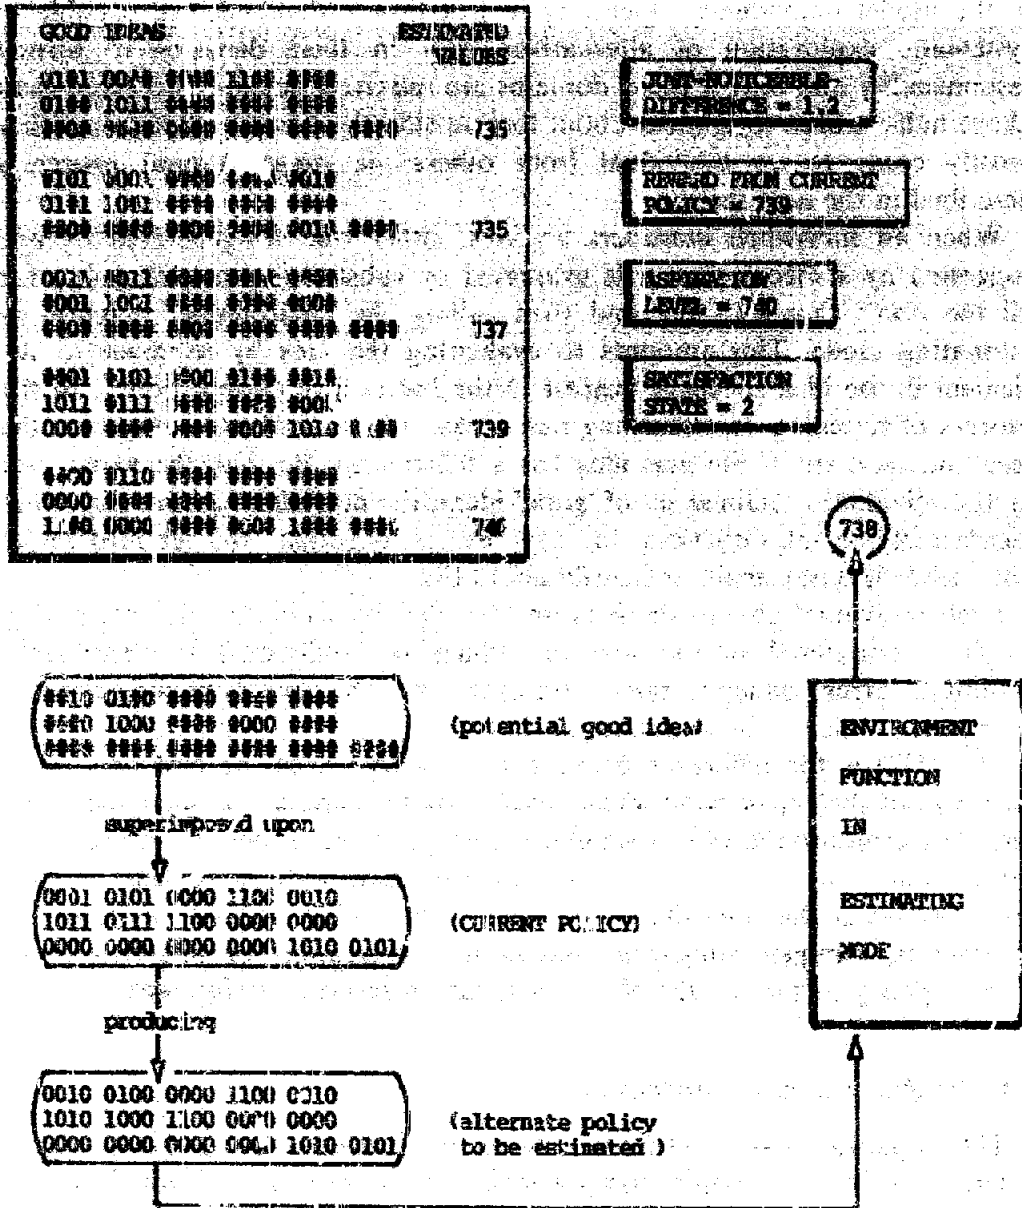


Fig 1. Basic elements that constitute a typical model individual in the situation shown a potential good idea is being considered. Its estimated value will result in it being placed fourth in the list of good ideas. The current fifth place entry will be dropped.

rather intensively. They do so both in isolation and when they explore how they would respond to policy changes suggested by others in meetings. Responsibilities may also enter into the choice procedures described below. A given locus may have one, many, or no responsible individuals.

The regular meetings of the organization occur in every time period and bring together fixed sets of participants over open agenda. A typical example

of a regular meeting process assembles meetings of individuals whose responsibilities are strongly interdependent. Each individual in such a meeting proposes a change in policy drawn from his or her list of good ideas. After the 'talk' interaction process, as explained below, the organization's choice procedures are invoked and produce a change in policy or a continuation of the status quo. Then the next meeting scheduled for the period takes place. This example is consistent with the models reported here. For a different type of investigation, quite different regular meeting processes can be readily defined by altering the regular meeting subroutine.

The *special meeting* subroutine provides a device for holding meetings over a tabled proposal that responds to particular problems or non-recurring conditions the organization may confront. An example might be a meeting called to bring together individuals for a discussion of a policy alternative with effects that cut across the boundaries implied by the regular meeting structure.

The *talk* interaction processes of a model consist of two phases. The first is the simple transmission of ideas to others. The second is determining and reporting what responses an individual would make in his or her own domain to ideas presented by others. The latter is analogous to answering the question 'If that idea were to become policy, what changes would I want to make in the loci for which I am responsible?' Since the result is also an idea, there can be responses to responses to suggested ideas, and so on. A limit of about four iterations is usually maintained. The organization may be provided with a theory of the relations between loci. The theory may be vague, or erroneous, or accurate. It is used by individuals in searching for responses to suggestions by others since they need to make some judgment of what implications a suggestion has for their domain of responsibility.

The *choice* procedures of a model specify for each locus the decision rule that must be used to make changes at that locus, a precedence structure that resolves conflicts over decision rules in multi-locus decisions, and a set of individuals who must be present for choices that affect a given locus. A typical model might have majority rule decisions for all loci, with responsible individuals required to be present. An alternative might give veto power over changes at some loci to particular individuals, or some loci might be subject to unilateral changes made by individuals responsible for them. Various combinations of unilateral, majority and veto choice procedures can be made by changing initial data for a model run. More complex changes can be made by altering the choice subroutine.

To gain a sense of how policy, environment, individual structure, and organizational structure are combined in a typical model it may help to consider the cycle of events experienced by an individual in one time period, as depicted in table 1. In an ordinary single cycle of model operations an individual receives the rewards generated for him or her by the policy in

Table 1

Typical model cycle

Begin period t .

For every individual i :

Update i 's reward from policy, aspiration level, just-noticeable-difference, and satisfaction state.

If current reward from policy > value estimated for policy alternative believed to be best, then update believed best.

Generate and consider one-locus variants of current good ideas; identify and make any preferred policy changes for which i has unilateral authority.

Generate and consider the opposite of current policy on loci of responsibility.

For every regularly scheduled meeting:

Assemble participating individuals.

Execute idea exchange processes (TALK routines).

Execute formal choice processes, changing policy when decision rule conditions are met.

For each special meeting called this period, if any:

Assemble participating individuals.

CONTINUE: If feasible agenda can be generated, then if any participant prefers it to status quo, then execute formal choice procedures.

If meeting has produced a policy change, then return to **CONTINUE**.

Collect statistical information and complete period t .

force at the completion of the previous period and updates aspiration, satisfaction and JND. Then he or she compares the current policy on loci of responsibility to the current idea of best policy on those loci. If the current policy is better, then the idea is adjusted. The individual estimates in isolation the value of the new ideas resulting from reversing one locus of responsibility at a time in current good ideas. If this generates an improvement it is added to the list of good ideas in place of the list's current minimum. If there are any changes that the individual thinks would be improvements on loci of responsibility and which the individual has authority to make unilaterally, they are made. The individual then attends those meetings that are on his or her regular schedule and special meetings to which he or she is invited, if any. On completion of whatever policy changes the meetings may engender, reports are generated and the period comes to a close.

3. The transport problem environment

In using the modelling system, one specifies not only a model organization, but also a model environment to which it adapts. The range of possible environment functions is large, but much of the possible variation can be generated by manipulation of one or more of three fundamental factors: the

pattern of interactions among the sixty-four policy loci, the nature of temporal variation in an individual's real or estimated rewards from a given policy, and the nature of interindividual differences in real or estimated rewards from a given policy. For example, the simplest possible environment would respond to each policy element equally and unconditionally, would give each individual accurate and temporally invariant estimates of the returns to be expected with any particular policy, and would give every individual in the organization the same real value for a particular policy. In this simplest linear case each policy locus has a setting that is preferable for all organizational members, and that setting is the same no matter what the setting of other loci. This environment presents a perfectly decomposable problem [Simon (1962)] to the model organization and can be solved by a moderate number of individuals in a single period.²

The linear environment without noise, change, or internal conflict is a case with minimal intrinsic interest. Model organizations have been run in numerous more interesting environments corresponding to highly non-linear integer programming problems and a program of extensive sampling of possible environments is now beginning. For the present paper, however, one particular type of environment has been studied in more detail. It is a classic family of transportation problems in 15 variables subject to 8 equality constraints. Table 2 shows the most distinguished member of the family, the

Table 2
Dantzig's original transportation example.

| | Warehouse | | | | | Factory capacity |
|----------------------|--|-----|-----|-----|-----|------------------|
| | 1 | 2 | 3 | 4 | 5 | |
| | <i>Costs of shipping/unit</i> | | | | | |
| Factory 1 | 9 | 20 | 18 | 12 | 25 | 500 |
| Factory 2 | 6 | 16 | 14 | 18 | 25 | 750 |
| Factory 3 | 27 | 18 | 15 | 10 | 9 | 250 |
| Warehouse capacities | 300 | 300 | 300 | 300 | 300 | 1500 |
| | <i>Optimal solution</i> | | | | | <i>Total</i> |
| Factory 1 | 150 | 0 | 0 | 300 | 50 | 500 |
| Factory 2 | 150 | 300 | 300 | 0 | 0 | 750 |
| Factory 3 | 0 | 0 | 0 | 0 | 250 | 250 |
| Total | 300 | 300 | 300 | 300 | 300 | 1500 |
| | Value of optimal feasible solution 18,350 | | | | | |
| | Value of pessimal feasible solution 30,950 | | | | | |

²A very small organization takes a little longer since responsibility for many loci taxes the limited attention of the individuals.

opening example of Dantzig's *Linear Programming and Extensions*. Dantzig's specification has been augmented by restricting the variables to integer values ranging from 0 to 375 in increments of 25.

The cost matrix and constraints correspond to the problem faced by an organization that must ship 25-case truckloads of products from three sources (factor²) to five destinations (warehouses³) without idling or exceeding the fixed capacity of the sources or the destinations. A feasible policy is any pattern of shipments that violates none of the constraints. An optimal policy is a feasible policy that minimizes total cost (shipping charges). A pessimal policy is a feasible policy that maximizes total cost.

The family of transportation problems related to the example of table 2 has been sampled by taking random integers from the interval [0, 25] for cost matrix entries. Ten new problems were constructed in this fashion and each organizational model in the discussion below was run in all ten of these model environments. (See appendix A.) None of the ten environments was used for testing during model development.

The integerized transport problem family makes a very challenging test for an organizational model. Feasible solutions are an extremely sparse subset of all policies.³ Policies generated merely by chance are therefore nearly certain to be infeasible. Moreover, feasible policies are not near neighbours. Since each of fifteen individuals in the model is responsible for one variable (four loci that are a binary code for the integers 0 to 15), any variation of a single element of a feasible policy produces an infeasible one for the models reported here. At least four loci that are responsibilities of four different individuals must be changed to transform one feasible policy into another. As many as ten loci might have to be reversed simultaneously. This means that discovering and implementing policy changes once a feasible policy has been found is cognitively difficult and requires the cooperation of several organizational members. When the interests of individual members are not identical the stage is set for what amounts to 'political' conflict over policy options. In essence, the integerized transport problems present a complex pattern that is not perfectly decomposable. They therefore provide a fundamental challenge to the *organization* of a set of boundedly rational individuals.

4. The basic organizational model

All the models run on these transport problems have most of their features

³A weak upper bound on the number of feasible policies can be derived by observing that there are only 91 ways a given warehouse can order 12 truckloads from three suppliers. If we ignore the factory capacity constraints we then have a number of feasible policies $\leq 91^5 \leq 6 \times 10^9$. The fifteen variables will use 60 policy elements, giving $2^{60} > 10^{18}$ total possible policies. Therefore the density of feasible policies is less — very substantially less, since we have ignored factory constraints — than one per hundred million.

in common. They are a selection of particular organizations from among the multitude that can be constructed with the modelling system described above. A number of features of the modelling system that might have introduced further realistic complexities were not used in these models, so as to simplify presentation and interpretation. The major features of the basic organizational model to be studied here may be grouped for description into six categories. In addition, we will examine the performance of two model variants, each of which is created by perturbing one key feature of the basic model.

4.1. Participants and their roles

All three models to be considered have sixteen members. Each member is responsible for four policy elements, which means that the portion of overall policy with which a particular member is most concerned can be in sixteen distinct states. Fifteen of the members have what might be called direct policy responsibilities. Each of them must consider what is the best level of ordering for a particular sink ('warehouse') from a particular source ('factory'). The combination of three sources and five sinks gives fifteen of these jobs. All of these fifteen individuals pursue the identical objective in the present model; for an alternative, see Cohen (1982). Possible order levels range from zero to fifteen units, where each unit corresponds to 25 'cases' in the original Dantzig specification. The sixteenth member has a managerial role. The state of the final four policy elements is used by the manager to control the incentive structure under which the other fifteen members operate. Each time policy stabilizes for more than a period with some factory constraints being violated, the manager can increase the (linear) penalty for such violations. Thus the manager responds simply to an observable condition: are the factory constraints satisfied? A precise definition of the environment faced by each individual is given in appendix B.

4.2. Routine meetings

The organizations all have a routine meeting schedule. In every time period, in every warehouse, the three individuals who place orders for that warehouse have a meeting in which each individual mentions the ordering policies currently held to be most attractive. The three consider each others' ideas and try to determine the implications of the implied changes in policy. They generate new ideas that fuse ideas mentioned by others and their own preferred responses to those ideas. At the conclusion of a meeting formal choice procedures may lead to changes in existing policy for their warehouse that have the support of at least two of the three members. In one of the model variants these three individuals may not all be from the same warehouse.

4.3. *Special meetings*

In addition, there is one meeting each period called by the manager, who selects from each of two randomly chosen warehouses two individuals who may be interested to discuss possible ordering patterns affecting both warehouses. These five try to find a reallocation of two single unit orders that seems to three or more of those attending to be an improvement on the current policy. If they make such a change, they consider whether a further similar change would be desirable. If so, it is made. If not, the meeting ends.

4.4. *Method of changing policies*

In the reported runs, majority rule voting was used in meetings to determine whether changes would be made. Individuals responsible for loci under discussion were required to participate in the choice. The only changes that could be made unilaterally were in the loci controlled by the manager, who could alter the incentives of other members single-handedly under the conditions described earlier.

4.5. *Theory of the environment*

In all runs, all individuals were provided with the identical rough theory of the organization's environment. A 64×64 incidence matrix represented the fact that orders for a given warehouse from a given factory were somehow interdependent with other orders for that warehouse and with other orders from that factory. It conveyed nothing about the nature of the interdependence. Thus any individual's efforts to determine how to respond to an idea suggested by another had only very weak theoretical guidance that identified one subset containing about one half of the loci as being of greater potential relevance.

4.6. *Initial policies and ideas*

All runs began with a policy of no orders. This infeasible policy has a disastrous value since violations of warehouse capacity constraints were always subjected to a severe quadratic penalty. All individuals began each run having no ideas about what to do other than the initial disastrous no-order policy.

The large scale structure of the basic organizational model is displayed in fig. 2. It shows fifteen member individuals in groups of three at each of the five warehouses and the sixteenth serving as manager. Each individual is shown adjacent to the set of four loci for which that member is responsible. Sixty loci control the patterns of shipping that in turn generate the organization's actual cost experience. The other four are used by the

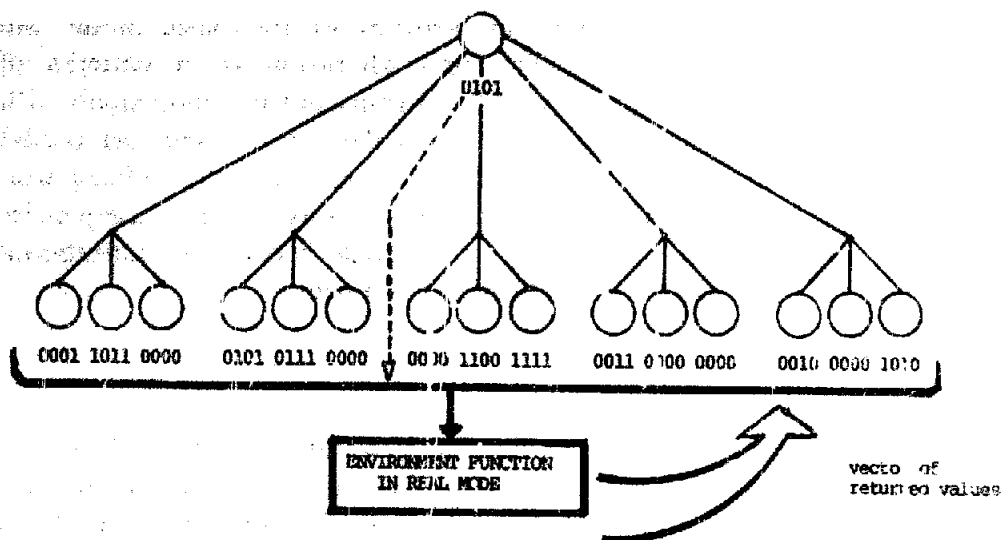


Fig. 2. Structure of the basic model. Each individual responsible for four loci, binary encodings of integers. Three individuals per warehouse at five warehouses, each determining this period's order for a warehouse from a factory. The loci of the sixteenth individual, the manager, control incentives of the others. The policy reaching the environment produces a vector of sixteen values. The underlying functions are in appendix B.

manager to adjust internal incentives. As each locus is a binary variable, there are 2^{60} possible patterns of shipment ordering. Organizational adaptation occurs as a result of the structured interaction of the search and choice processes of the member individuals. We want to assess the contribution of that structure to the quality of the organization's adaptive performance in this very complex environment.

5. The two variant models

The results to be reported below are derived from the performance of three models as each faced all ten variations on the Dantzig environment. These three are the basic model, described above, and two variants which may be labelled 'No-Talk' and 'Random-Meetings'.

In the No-Talk model, the talk processes described above are partially disabled. Individuals still learn of ideas preferred by others, but they do not go on to consider, or pass on to others, what they would do — or prefer to do — if such ideas were to be implemented.⁴ Since the No-Talk variant is similar in all other respects to the basic model, it can be run in the same test model environments in order to assess the specific contribution of the short-circuited interaction processes to overall organizational performance.

⁴Mechanically this is accomplished simply by inserting a REFL R/S instruction in the basic model's TALK subroutine. (See IPPS discussion paper no. 151 for details.)

In the Random-Meetings variant all features of the basic model are present, including talk processes. The only alteration is in patterns of attendance at regular meetings. Instead of consisting of three individuals with responsibilities at the same warehouse, each of the five meetings per period consists of three individuals drawn randomly from among the fifteen who have direct policy responsibilities. This permits us to assess the consequences of a simple form of misorganization, a mismatch between organizational structure and the structure of the problem environment.

6. Results

These extensive preliminaries now allow us to report the results of running each of the three model organizations in the ten environments that are variants on the Dantzig transport problem. Table 3 shows in its left-hand column the remarkable adaptive power of the basic organizational model. In eight of the ten environments the model was able to find optimal policies. In the remaining two the policy reached within thirty-five time periods was close to the optimum by two separate measures: percent of range of feasible values, and number of constraints violated by policy.

In the central column of the table one sees the dramatic difference in performance that follows from short-circuiting the TALK subroutine of the basic organizational model. The effect on the quality of organizational adaptation is disastrous. No optimal solutions were obtained. Indeed, no feasible solutions were found. No policy at the end of thirty-five periods in the modified model was within sight of the worst performance by the basic model.

In the right-hand column of the table it is evident that misorganization of the interaction processes damages organizational adaptation. Not just any structure to channel the processes embodied in the TALK subroutine will do. Randomization of regular meeting attendance patterns results in policies substantially worse than those obtained by the basic model in every one of the ten environments although TALK interaction was not disabled.

Although both variants were far behind the performance of the basic model, the Random-Meetings model seems to have been less completely hobbled than the No-Talk variant, and the reasons are instructive. In two environments the Random-Meetings model did attain feasible, though not optimal, policy. In two others the feasible region was closely approached. The No-Talk variant never approached feasibility and failed to do so in a very striking way: nine of the ten runs have warehouse constraints fully satisfied while factory constraints are virtually ignored. In No-Talk, warehouse meetings occurred routinely, and it was usually possible for at least one of the three individuals to think of a policy for the warehouse that satisfied its constraint. Without the TALK subroutine, however, the groups were unable

Table 3.
Performance of three model variants.

| Experiment | Basic model | | | | No-Talk variant | | | | Random meetings | | | |
|------------|-----------------------|---|---|---|-----------------------|---|---|---|-----------------------|---|---|---|
| | Optimal final policy? | Suboptimal performance and penalties as % of feasible range | Number of constraint violations (Column, row) | Number of constraint violations (Column, row) | Optimal final policy? | Suboptimal performance and penalties as % of feasible range | Number of constraint violations (Column, row) | Number of constraint violations (Column, row) | Optimal final policy? | Suboptimal performance and penalties as % of feasible range | Number of constraint violations (Column, row) | Number of constraint violations (Column, row) |
| 1 | Yes | 0% | 0,0 | 0,0 | No | — | 2,12 | 2,12 | No | 1.7% | 0,0 | 0,0 |
| 2 | Yes | 0% | 0,0 | 0,0 | No | — | 0,60 | 0,60 | No | — | 3,33 | 3,33 |
| 3 | No | 2.5% | 0,0 | 0,0 | No | — | 0,18 | 0,18 | No | — | 5,5 | 5,5 |
| 4 | Yes | 0% | 0,0 | 0,0 | No | — | 0,18 | 0,18 | No | — | 2,26 | 2,26 |
| 5 | No | 3.5% | 0,0 | 0,0 | No | — | 3,32 | 3,32 | No | 19.8% | 0,0 | 0,0 |
| 6 | Yes | 0% | 0,0 | 0,0 | No | — | 0,52 | 0,52 | No | — | 12,38 | 12,38 |
| 7 | Yes | 0% | 0,0 | 0,0 | No | — | 0,20 | 0,20 | No | — | 2,10 | 2,10 |
| 8 | Yes | 0% | 0,0 | 0,0 | No | — | 0,52 | 0,52 | No | — | 1,29 | 1,29 |
| 9 | Yes | 0% | 0,0 | 0,0 | No | — | 0,46 | 0,46 | No | — | 1,1 | 1,1 |
| 10 | Yes | 0% | 0,0 | 0,0 | No | — | 0,44 | 0,44 | No | — | 2,2 | 2,2 |

to find policies that did this without disastrous effects on factory constraints. In the resulting policy instability, penalties for factory constraint violations tended to remain at low levels, and this exacerbated the problems. In the Random-Meetings model six of the cases have factory violation levels better than the best case — No-Talk. In this small organization, chance meetings of three individuals will include a pair from a common warehouse about forty percent of the time, and a pair from a common factory will be present in about seventy percent of the meetings. In these cases TALK has some chance to work — although with possible distortion introduced by the third party present. Even though its opportunities are unreliable and possibly distorted, the interaction process is powerful enough to make substantial contributions in random meeting. This does not occur in the No-Talk variant.

These results establish three major conclusions: (1) an appropriate organization of boundedly rational individuals is capable of optimal or near optimal adaptive performance in an environment of substantial complexity, (2) high quality adaptive performance can be derived principally from the structured interaction of the weak individuals who compose the organization, and (3) an inappropriate structure for that interaction will achieve markedly inferior adaptation.

Strictly speaking these conclusions pertain, of course, only to the model world. Their significance beyond that depends on one's confidence in the relevance of the model world to the real one. That in turn must rest chiefly on the detailed efforts made to incorporate widely observed characteristics of organizations and their decision makers in the model.⁵

The conclusions about the model world alone, however, are sufficient to establish that the decision making approach to organizations can account for the occurrence of high quality organizational adaptation to very complex environments. It can do so in a fashion which makes adaptation a property of the organization rather than of its individual members, and which permits successful adaptation without guaranteeing it. This is a partial fulfillment of the promise of Simon's (1964) profound treatment of organizational conflict and goals.

7. Interpretation

There are four points that should be made in interpreting these model results. The first of these is that *a central role is played by search processes when an organization confronts a difficult environment*. The point is hardly new, especially when considering problem solving by individuals [Simon

⁵ That confidence may be augmented by the results of work currently under way which will compare predictions from model variants to performance by human groups operating in corresponding conditions, but these studies are not yet complete.

(1955 and 1956)]. It is worth stressing, however, that its formal support has been strengthened by recent developments in the theory of computational complexity and that the assertion is, in the present case, made not about the search of individuals, but about the search of organizations.

Since Simon's early articulation of bounded rationality arguments, there has emerged a very substantial literature on the restricted ability of individuals to consider alternatives [Slovic et al. (1977)] or to handle probability concepts. That empirical literature has now been reinforced by theoretical analysis of what makes problems computationally complex. Beginning with Cook's (1971) theorem theoretical computer scientists have established that a daunting range of problems of economic significance are NP-complete. The conjecture has become general that such problems can require the consideration of an overwhelming number of alternatives and are therefore beyond the reach of any algorithm running in time bounded by a polynomial function of the number of variables in the problem [Garey and Johnson (1979)]. In this situation, search heuristics have become a topic of major research interest since exact optimal solutions to problems of theoretical and economic significance are unlikely to be obtained [Karp (1976), Weiner (1975), Rosencrantz et al. (1974), Maffioli (1979)]. Computer science has thus come up against the same kinds of boundaries that limit the problem-solving activities of human beings, and has had recourse to the same strategy in response. Combinatorial complexity makes the design of search processes a crucially important activity.

Real organizations obviously face problems at least as complex as the simplified models of those problems formulated in operations research and computer science, so it is entirely plausible that the quality of search be a major determinant of the quality of organizational decision making. The corresponding result has been obtained from the model. The organizational search processes taking place in the TALK routine that communicates ideas and counterideas among individuals have been shown to be capable of generating optimal and near-optimal results to problems of extreme difficulty.⁶ It has also been shown that the organizational search processes available when the interindividual activity is blocked, are incapable of high quality adaptation. The pair of results gives rise to the second major point to be made: *the quality of the model's search is, to a substantial degree, a property of the organization and not merely the summation of properties of the individuals who compose the organization.* In this model — and, so far as it is valid, in the world — interactions among individuals constitute a principal source of phenomena that are distinctively organizational. Those interactions

⁶Technically the model organization is facing a Zero-One Integer Programming problem in sixty variables. The Zero-One problem is NP-complete (Berthel and Erving 1970, Karp 1972). We know that this one can be reduced to a transport problem for which an optimal method is available, but the model does not 'know' that.

are a major source of the augmented decision making power which the organization can display.

One brings us to the third point to be made: organizational structure is, among its other functions, a search heuristic. The quality of organizational search is fundamentally dependent critically on the division of responsibility and attention among the organization's members, on the patterns of interaction among members induced by regular and special meetings, and on the pattern of diverse incentives which shape the kinds of new policies to which the individuals and groups are predisposed. In effect, the structure of the organization, its formal and informal structuring of interactions, and its biasing incentives, are a concrete heuristic for searching the vast space of logically possible policies.⁸ And, for the model, they are necessary for success.

The fourth and final point is that the principal heuristic employed by the model (viz. the structuring of organizational processes to produce powerfully effective search) is strongly parallel. That is, it derives its power from the interaction of multiple concurrent processes.⁹ The No-Talk and Random-Meeting cases show the effects of blocking or misstructuring the interactions among the searches for better policy being conducted by the organization's members. These gains from organization and interaction are not merely speed improvements derived from doing independent parts of a job simultaneously. In the No-Talk variant those advantages are still being realized, but the results achieved are nonetheless vastly inferior to those of the basic model. The basic model's performance rests on the logical and temporal appropriateness of the interactions to which it naturally gives rise.

The striking feature of the strongly parallel heuristic of the model is its ability to get powerful search performance out of weak parts. This characteristic is also exhibited in some other areas where parallel processes have been shown to dramatically outperform nonparallel alternatives. For example, Fahlman (1979) has shown that set intersections may be found in one or a few cycles by his parallel information network no matter what size the sets. Typical serial processes will take a time proportional to set size. Thus Fahlman's system can determine almost immediately, as can most

⁸A companion paper, 'Conflict and Complexity' explores this aspect in more detail [Cohen (1982)]. It shows that there are circumstances in which the divergence of incentives may protect an organization against misjudgments of a hard-to-understand environment.

⁹Other functions are clearly served by structure. See, e.g., Cremer (1980), or Thompson (1967).

⁹It is important — and, unfortunately, not common — to distinguish this kind of parallelism from merely simultaneous but nearly independent processes. Kornfeld (1979, 1981) and Inai et al. (1979) are the only instances I have found that recognize the difference clearly. Kornfeld calls strongly interacting parallelism 'combinatorial implosion'; Inai et al. use 'acceleration effect'. The execution of the computer program embodying the model is, of course, serial by virtue of the nature of most computer systems currently available. The program itself, however, simulates the strong parallelism of organizational processes.

humans, that it knows of no insects larger than a horse.¹⁰ A serial system not designed in advance for just this type of question will bog down in the large size of the sets. A second example is provided by the work of John Holland (1975) on heuristic exploration of high-dimensionality non-linear functions. He and his colleagues have shown that a population of about a hundred 'individuals' interacting according to rules borrowed from genetics can find global optima in such conditions much more reliably than standard hill climbing techniques using comparable amounts of computer time.¹¹ A third example is provided by the remarkable improvements in computation promised by highly parallel algorithms exploiting the massive simultaneity that will be possible with Very Large Scale Integrated architectures. Mead and Conway (1980, pp. 264-265) suggest that 'analogies with human structures may help to suggest the kinds of behavior we might achieve in computational structures . . . [Parallelism] is widely exhibited in human organizations . . . The design of computers and of algorithms has yet to show the ingenuity reflected in human organizations'.

If the essential problem of organizational adaptation is coping successfully with the complexity of the environment, and if we believe the individuals in the organization are generally not capable of solving its whole problem alone, this property of generating powerful performance from interactions of weak components will have to be present for high quality adaptation to occur. The results reported above show that interacting parallel processes broadly consistent with modern empirical studies of organizations and their decision makers are sufficient to produce impressive adaptation. These processes allow the basic model to meet a stringent and essential performance test that previous accounts of organizational adaptation have not faced and probably could not meet. As a result, fundamental questions about the limits and sources of organizational rationality may now be subjected to a more rigorous and revealing theoretical examination.

¹⁰In an organizational setting, an equivalent question would be whether there is a known policy that satisfies several given constraints.

¹¹In fact, this demonstration by Holland provided a fundamental impetus for the construction of precursors of the model organizations reported here.

Appendix A: Dantzig problem variants

Table A.1

| Problem | Cost matrix | Original solution (in 25 unit truckloads) | Best value/worst |
|---------|----------------|--|------------------|
| 1 | 21 23 5 24 12 | 8 0 12 0 0 | 670/912 |
| | 15 10 5 16 5 | 4 12 0 2 12 | |
| | 27 6 8 13 7 | 0 0 0 10 0 | |
| 2 | 11 17 13 12 5 | 0 0 0 8 12 | 765/1026 |
| | 9 16 24 18 14 | 12 2 12 4 0 | |
| | 8 11 22 19 12 | 0 10 0 0 0 | |
| 3 | 9 20 10 15 15 | 6 0 12 0 2 | 622/1098 |
| | 6 8 23 8 17 | 6 12 0 2 0 | |
| | 21 17 18 24 19 | 0 0 0 0 10 | |
| 4 | 7 17 11 17 23 | 8 0 12 0 0 | 656/1124 |
| | 12 17 24 6 12 | 4 2 0 12 12 | |
| | 14 5 22 15 24 | 0 10 0 0 0 | |
| 5 | 20 13 10 11 21 | 0 0 8 12 0 | 752/1236 |
| | 18 16 22 18 8 | 6 12 0 0 12 | |
| | 16 24 12 17 10 | 6 0 4 0 0 | |
| 6 | 18 17 5 15 22 | 0 2 12 0 6 | 858/1200 |
| | 18 23 6 14 23 | 12 0 0 12 6 | |
| | 12 11 24 7 19 | 0 10 0 0 0 | |
| 7 | 24 13 5 13 7 | 0 0 12 0 8 | 476/1142 |
| | 8 12 18 11 21 | 12 6 0 12 0 | |
| | 12 6 13 21 6 | 0 6 0 0 4 | |
| 8 | 15 20 16 21 6 | 0 0 8 0 12 | 742/1078 |
| | 6 21 24 23 13 | 12 6 0 12 0 | |
| | 10 6 8 11 6 | 0 6 4 0 0 | |
| 9 | 18 9 15 5 17 | 0 0 8 12 0 | 570/914 |
| | 12 6 15 9 17 | 12 12 4 0 2 | |
| | 15 5 14 20 8 | 0 0 0 0 10 | |
| 10 | 6 15 20 15 19 | 12 0 0 8 0 | 652/1082 |
| | 22 9 16 17 11 | 0 2 12 4 12 | |
| | 16 5 12 14 13 | 0 10 0 0 0 | |

Appendix B

Definitions

x_{ij} \equiv order level, $0 \leq x_{ij} \leq 15$.

x_{ij} may be interpreted as 'truckloads'. $1 \leq i \leq 3$, factories, $1 \leq j \leq 5$, warehouses,

v_{ab} \equiv value for policy returned to individual (a, b),

c_{ij} \equiv cost of shipping one ordered truckload from i to j ,

r_i \equiv total of orders from factory $i = \sum_{j=1}^5 x_{ij}$,

$p = |r_1 - 20| + |r_2 - 30| + |r_3 - 10| \equiv$ sum of factory capacity violations,

$w_j = \sum_{i=1}^3 x_{ij} - 12 \equiv$ capacity violation at warehouse j .

For directly responsible individual (a, b) ordering from factory a for warehouse b ,

$$y_{ab} = \sum_{i=1}^3 \sum_{j=1}^5 x_{ij} c_{ij} + (0.55p)(c(K+1) - c(K)) + \sum_{j=1}^5 (5w_j)^2.$$

The individual acts to minimize this function of total shipping cost at all warehouses, factory capacity violations, and all warehouse capacity violations, $c(K)$ is the K th smallest unit cost and K is increased by the manager each time organizational policy has been unchanged for three consecutive periods.

For the manager

$$y_m = 0.0001 \sum_{i=1}^3 \sum_{j=1}^5 x_{ij} c_{ij} + p.$$

The scaling factor (0.0001) establishes a lexicographic preference for reducing p , factory constraint violations, and thereafter for alternatives that minimize total shipping cost.

References

- Aldrich, H., 1979, *Organizations and environments* (Prentice-Hall, Englewood Cliffs, NJ).
- Anderson, J., 1980, *Cognitive psychology and its implications* (Freeman, San Francisco, CA).
- Borosh, I. and L.B. Trevbig, 1976, Bounds on positive integral solutions of linear diophantine equations, *Proceedings of American Mathematical Society* 55, 299-304.
- Cohen, M., 1980, Documentation of an organizational modelling system, Discussion paper no. 151 (Institute of Public Policy Studies, University of Michigan, Ann Arbor, MI).
- Cohen, M., 1982, Conflict and complexity, Discussion paper no. 153, (Institute of Public Policy Studies, Ann Arbor, MI) revised.
- Cook, S.A., 1971, The complexity of theorem proving procedures, *Proceedings Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York).
- Greiner, J., 1980, A partial theory of the optimal organization of a bureaucracy, *Bell Journal of Economics*, Autumn, 683-693.
- Estes, W.K., 1978, *Handbook of learning and cognitive processes*, Vol. 6 (Erlbaum, NJ).
- Fahlman, S., 1979, *NETL: A system for representing and using real-world knowledge* (MIT Press, Cambridge, MA).

- Garey, M.R. and D.S. Johnson, 1979, *Computers and intractability* (W.H. Freeman, San Francisco, CA).
- Holland, John, 1975, *Adaptation in natural and artificial systems* (University of Michigan Press, Ann Arbor, MI).
- Imai, M., T. Yamada and T. Futamura, 1979, A parallel searching scheme for multi-processor systems and its application to combinatorial problems, *3rd International Joint Conference on Artificial Intelligence*, Aug.
- Karp, R., 1972, Reducibility among combinatorial problems, in: Miller and Thatcher, eds., *Complexity of computer computations* (Plenum Press, New York).
- Karp, R., 1976, The probabilistic analysis of some combinatorial search algorithms, in: J.D. Traub, ed., *Algorithmic and complexity: new directions and recent results* (Academic Press, New York).
- Kornfeld, W.A., 1979, Using parallel processing for problem solving, *Artificial Intelligence Lab memo no. 561*, Dec. (MIT, Cambridge, MA).
- Kornfeld, W.A., 1981, The use of parallelism to implement a heuristic approach, *Seventh International Joint Conference on Artificial Intelligence*, Aug.
- Maffioli, F., 1979, The complexity of combinatorial optimization algorithms and the challenge of heuristics, in: N. Christofides, ed., *Combinatorial optimization* (Wiley, New York) 107-130.
- Mead, C. and L. Conway, 1980, *Introduction to VLSI systems* (Addison-Wesley, Menlo Park, CA).
- Newell, A. and H.A. Simon, 1972, *Human problem solving* (Prentice-Hall, Englewood Cliffs, NJ).
- Pfeffer, J. and G.R. Salancik, 1974, *The external control of organizations* (Harper & Row, New York).
- Rosencrantz, D.J., R.E. Stearns and P.M. Lewis, 1974, Approximation algorithms for the travelling salesman problem, *Proceedings of the 15th IEEE Switching & Automata Theory Symposium*.
- Simon, H.A., 1955, A behavioral model of rational choice, *Quarterly Journal of Economics* 69, Feb.
- Simon, H.A., 1956, Rational choice: and the structure of the environment, *Psychological Review* 63, March.
- Simon, H.A., 1962, The architecture of complexity, *Proceedings of the American Philosophical Society* 106, Dec., 457-482.
- Simon, H.A., 1974, On the concept of organizational goal, *Admin. Sci. Quarterly* 9, June, 1-22.
- Simon, H.A., 1969, *Sciences of the artificial* (MIT Press, Cambridge, MA).
- Slovic, P., A. Fischhoff and S. Lichtenstein, 1977, Behavioral decision theory, *Ann. Rev. of Psych.* 1-39.
- Starbuck, W.H., 1973, Organizations and their environments, in: M. Dunnette, ed., *Handbook of industrial and organizational psychology* (Rand McNally, Chicago, IL).
- Thompson, J., 1967, *Organizations in action* (McGraw Hill, New York).
- Weiser, P., 1975, Heuristics, *Networks* 5, no. 1, Jan., 101-113.
- Winter, S., 1975, Optimization and evolution in the theory of the firm, in: R.H. Day, ed., *Adaptive economic models* (Academic Press, New York).