Carnegie Mellon University
Research Showcase

Department of Electrical and Computer Engineering

Carnegie Institute of Technology

1-1-1996

Cooperation schemes for autonomous agents

Sarosh Talukdar
Carnegie Mellon University

Carnegie Mellon University.Engineering Design Research Center.

Follow this and additional works at: http://repository.cmu.edu/ece

Recommended Citation
Talukdar, Sarosh and Carnegie Mellon University.Engineering Design Research Center., "Cooperation schemes for autonomous agents" (1996). *Department of Electrical and Computer Engineering.* Paper 201.
http://repository.cmu.edu/ece/201

This Technical Report is brought to you for free and open access by the Carnegie Institute of Technology at Research Showcase. It has been accepted for inclusion in Department of Electrical and Computer Engineering by an authorized administrator of Research Showcase. For more information, please contact research-showcase@andrew.cmu.edu.

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Cooperation Schemes for Autonomous Agents

Sarosh Talukdar, Lars Baerentzen,
Andrew Gove and Pedro de Souza

EDRC 18-56-96

COOPERATION SCHEMES FOR AUTONOMOUS AGENTS

Sarosh Talukdar

Lars Baerentzen

Andrew Gove

Pedro de Souza

Carnegie Mellon University
Pittsburgh, PA 15213

Copyright 1996 by Talukdar

COOPERATION SCHEMES FOR AUTONOMOUS AGENTS

Sarosh Talukdar Lars Baerentzen Andrew Gove Pedro de Souza

Engineering Design Research Center
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

Experiments over a variety of optimization problems have shown that convergence to good solutions is an emergent behavior of certain mixes of autonomous (unsupervised) agents in certain cooperative arrangements. What mixes and arrangements? What are the underlying mechanisms? What are the implications for organization design? This article provides some answers.

I. INTRODUCTION

None of the known algorithms for optimization and constraint satisfaction is without weaknesses-the rigorous algorithms tend to be too slow and cumbersome, the heuristics, too unreliable. Rather than seeking a new and better algorithm, we have been experimenting with ways by which available algorithms and other computer-based operators can cooperate, so together they can do what separately they might not. The result is a type of organization, called an asynchronous team (A-Team), that combines features from a number of systems, particularly, insect societies [1], cellular communities [2], genetic algorithms [3], blackboards [4], simulated annealing [5] and tabu search [6].

Definition: an A-Team is an evolving sequence of strongly cyclic data flows. A data flow is a directed hypergraph like those shown in Fig. 1. Nodes in a data flow are Venn diagrams representing complexes of overlapping memories, or more precisely, the objects these memories could contain. Arcs in a data flow represent agents that read from the memories at their tails and write to the memories at their heads. A data flow is strongly cyclic if all, or almost all, its arcs are in closed loops.

A-Teams are obtained by forming operators into autonomous agents and agents into strongly cyclic data flows that are implemented in networks of computers. Each A-Team is dedicated to one problem. Each data flow is a distinct scheme for solving instances of this problem. Results (trial-solutions) accumulate in the memories of the data flow (just as they do in blackboards) to form populations (like those in genetic algorithms). These populations are continually modified by two types of agents: construction agents that add members to the populations and destruction agents that eliminate members from the populations. The latter work from lists of solutions to be avoided (like the lists used in tabu searches).

The numbers of construction and destruction agents can be arbitrarily large and each agent, whatever its type, can be arbitrarily complex. Consequently, the problem-solving skills of a data flow can be arbitrarily apportioned between construction and destruction. (Other synthetic problem-solving systems invariably concentrate on one or the other. Hill climbing, for instance, concentrates on how to construct new and better solutions while simulated annealing, genetic algorithms and tabu search concentrate on how to destroy or reject weak solutions. Natural systems, however, often benefit from a more symmetric use of construction and destruction. The process of Lamelar bone growth [9], for instance, relies as much for its efficacy on cells that add bone material to surfaces where the stress is high, as it does on cells that remove bone material from surfaces where the stress is low.)

All the agents in every data flow of an A-Team are autonomous. An autonomous agent decides for itself what it is going to do and when (like the adult members of insect societies). Consequently, there can be no centralized control. But new autonomous agents can be easily added (there is no centralized control system to get in the way).

Agents cooperate by working on one another's results. Because the agents are autonomous, this cooperation is asynchronous (no agent can be forced to wait for results from another). Rather, all the agents can, if they so choose, work in parallel all the time. (Other

synthetic problem-solving systems often include precedence constraints to force at least a partial order on the activities of their computational modules. Traditional genetic algorithms, for instance, require destruction to cease while construction is in progress, and vice-versa.)

Since every agent decides for itself what, if anything, to do and when, if ever, to do it, and since every agent is unaware of its colleagues except for the results they produce, one might think that the agents would work at cross purposes. Surprisingly, this not always the case. Useful A-Teams have been developed for a wide variety of optimization and constraint satisfaction problems, including, nonlinear equation solving [7], [24], traveling salesman problems [14], high-rise building design [8], reconfigurable robot design [9], diagnosis of faults in electric networks [10], control of electric networks [11], [25] job-shop-scheduling [16], steel mill scheduling [17], paper mill scheduling [25], [26], train-scheduling [15], and constraint satisfaction [18]. Not only do the data flows of these A-Teams produce very good solutions, but they appear to be scale-effective, that is, the data flows can always be made to produce even better solutions by the addition of certain agents and memories.

Scale-effectiveness is rare in synthetic organizations. Hence the proverb, "too many cooks spoil the broth." In a scale-effective organization there can never be too many cooks, at least not from the perspective of broth-quality. More precisely, the problem of obtaining better solutions in a scale-effective organization reduces to one of finding which components to add. A non-scale-effective organization usually faces the much more difficult problem of finding which of its parts to eliminate or modify before additions can be of benefit.

II. A DESIGN PROTOCOL

For several years, we have been experimenting with a protocol for designing A-Teams. The steps of this protocol together with some experience-based guidelines for their implementation are given below and illustrated with an A-Team for the traveling salesman problem (TSP).

Step 1. Choose a problem.

The TSP is a prototypical combinatorial-optimization-problem with a variety of practical applications. In its most basic form, the TSP is: given m cities and their separations, find the shortest tour of the cities. (A tour is a closed path that goes through every city). The number of distinct tours grows so rapidly with m that it is impossible to conduct an exhaustive search for the shortest tour, even when there are as few as 30 cities. Practical instances the TSP often have hundreds and sometimes thousands of cities.

Step 2. Decompose the problem into subproblems.

The decomposition need not be hierarchic. The subproblems need not be disjoint, nor even distinct; they may have only small overlaps with the original problem and may even be more complex than it. But subproblems should be chosen so they can be connected by agents into closed loops, so the solutions of each subproblem can be used to construct solutions to the next subproblem in the loop.

For the TSP, we choose four subproblems: find an optimal tour, find good tours, find good partial tours, and find good 1-trees; where "good" means "containing many of the arcs of an optimal tour" and a 1-tree connects all the cities but does not always form a tour (the 1-trees form a superset of the tours).

Step 3. Assign one or more memories to each subproblem.

The purpose of each memory is to hold a population of trial-solutions to its subproblem. As in the case of genetic algorithms, larger populations lead to better solutions. But the marginal benefits fall off rapidly. Therefore, moderately sized populations can be expected to work as well as very large ones.

Even moderate populations can occupy a great deal of storage space, especially when each trial-solution is a complex object, such as a building-design. In such cases, the representation should be a carefully chosen compromise between compactness and clarity.

For the TSP, we choose m as the population size (m/2 and 2m seem to work equally well) and ordered lists of cities as the representation for complete and partial tours. With this representation, {Atlanta, Boston, Raleigh, Pittsburgh} means a tour that goes from Atlanta to Boston to Raleigh to Pittsburgh, and then back to Atlanta. The representation chosen for 1-trees is more complicated. Interested readers can find it in [14].

Step 4. Select a set of algorithms or operators for each subproblem.

The greater the range of skills of these algorithms, the better the solutions that will be found. The algorithms do not have to be uniform in size or coverage. Rather, some can be large, others small, some can be general, others specialized. From their experiences in developing commercial A-Teams for industrial clients, Murthy and his colleagues [26], [27]f recommend using all the best and most powerful algorithms that are available. Weak algorithms, such as crossover and mutation, when used exclusively, or powerful algorithms, such as a state-of-the-art-branch-and-bound-algorithm, when used alone, or any mix of low variety, will invariably do less well than a mix of greater variety and range.

For the TSP, rigorous algorithms that finish in reasonable (polynomial) amounts of time are unknown. However, there are hundreds of faster heuristics for generating and refining sub-optimal tours. We chose the sample shown in Fig. 2.

Step 5. Form each algorithm into an autonomous agent.

Think of an agent as having three components: an operator, a communication system and a control system, the communication system connects the agent to its designated input and output memories. The control system consists of selectors and schedulers. The former choose objects from the input memory for the operator to work on; the latter determine when the operator will work and which of the available computers it will use.

> Definition: an agent is autonomous if its control system is completely self-contained, that is, if it accepts no selection or scheduling instructions from other agents.

The key to effective cooperation among autonomous agents is in the design of their selection strategies. Where the quality of the solution can be measured by a single attribute (such as tour length for the TSP), a very simple selection strategy seems to work quite well. This strategy is a mirror image of the solution- rejection-strategy used in simulated annealing, specifically: select solutions randomly with a bias towards the better solutions. Murthy [9] has devised a variant of this strategy for cases where solution-quality is best measured by a vector of conflicting attributes. Specifically: compare a vector representing the estimated needs of the solution to a vector representing the estimated capabilities of the agent; arrange for the probability of selection to increase as the magnitude of the angle between these vectors decreases.'

While scheduling strategies are undoubtedly important, we have yet to investigate their effects. For the TSP and all the other cases we have studied, we have used only one very simple strategy: allow each agent to run continuously, or as close to continuously as the available computers will permit.

Step 6. Form autonomous destroyers.

Agents can be of two types: constructors (that add trial-solutions to their output-memories) and destroyers (that erase trial-solutions from their output-memories). The agents produced in the previous step are constructors. Their actions must be balanced by destroyers or the memories would soon become clogged with trial-solutions.

Destroyers can serve two additional functions. First, by erasing any solutions that fall in parts of their output spaces (the sets of all the solutions that can be stored in their output-memories) destroyers can make those parts almost inaccessible to the constructors. (For a region to be truly inaccessible, results would have to be prevented from falling in it, not merely erased after falling in it.) Second, destroyers can terminate undesirable patterns of constructor behavior, such as (repeating sequences of solutions, by recognizing and erasing them.

For the TSP, we used destroyers to erase tours selected randomly with a strong bias for the longest tours (just as in the solution-rejection-strategy of simulated annealing). Destroyers scheduled their activities so there were always a few open slots for the constructors to fill.

Step 7. Form the agents and memories into a strongly cyclic data flow.
Some of the data flows developed for the TSP are shown in Fig. 1.

Step 8. Test and modify the data flow by implementing it in a network of computers.
Seed the memories with initial populations of trial-solutions, activate the agents and monitor the changes in the solution-populations. If these changes are overly slow in converging, that is, if at least one complete solution of acceptable quality is overly slow in appearing, then repeat from step 4. If convergence is still too slow, repeat from step 2.

Thus, a sequence of data flows is produced by iterating through two nested loops. The outer loop, consisting of steps 2-8, changes the decomposition (the set of memories). The inner loop, consisting of steps 4-8, stocks the data flow with algorithms. When the supply of these algorithms is exhausted, or when the marginal benefit of adding new algorithms grows small, the outer loop is reinvoked in the quest for a new and more powerful decomposition. In iterating through these loops, proceed under the assumption that strongly cyclic data flows can easily be made scale-effective; it is best to begin with a small data flow, then add agents and memories until a data flow with adequate performance is obtained.

The results for the TSP are shown in Figs. 3 and 4. Notice the effects of scale on solution-quality, especially, how quality improves as agents and memories are added. Notice also that the speed with which the results were produced often increased with increases in scale, even though all the agents and memories were made to share a single computer (Fig.3). When more computers were made available (Fig. 4), the larger data flows invariably finished faster than the smaller ones. Specifically, when four or more computers were made available to each data-flow, 2 (d), the largest data flow, not only produced

the best solutions but produced them faster than any other data flow. Thus, for the A-Team for $TSPs_t$ solution-quality and speed appear to be commensurate--both improve as certain agents and memories are added.

Observed Behavior

The protocol has been tested over a wide variety of optimization problems [7-11, 14-18, 24-27]. The results display similarities along six dimensions. These dimensions and the associated problem-independent behaviors are:

1. Diversity: solution-quality increases with the range of skills of the construction agents.
2. Scale: there is little, if any, penalty for an excess of construction skills. Rather, scale-effectiveness seems to be commonplace; solution quality'can invariably be improved by the addition of construction and destruction agents.
3. Expansion: adding autonomous agents to strongly cyclic data flows is relatively easy, regardless of whether the agents are large or small, general or specialized.
4. Duality: adept destruction can compensate for inept construction, and vice-versa.
5. Population size: solution-quality benefits from increasing the sizes of solution-populations, but these benefits are prone to saturation.
6. Parallelism: solution-speed improves as computers are added until there are enough computers for all the agents to work in parallel all the time. Often, the speed-up is near-linear.

III. COOPERATION IN STRONGLY CYCLIC DATA FLOWS

What mechanisms are responsible for the problem-independent behaviors of A-Teams? What are the underlying phenomena and causal relations? We will tackle these questions with the aid of a device, called a CDM (constant drift memory) which, it can be argued, is a good model of A-Teams. The argument in outline is:

- all forms of cooperation can be modeled by data flows;
- in any data flow, all the behaviors of interest occur in only one memory;
- this, and all the other memories in any strongly cyclic data flow, can be accurately modeled by devices called cyclic memories;
- most, if not all, cyclic memories can be modeled by a particularly simple type of cyclic

memory called a CDM.;

The technical apparatus needed to make this argument is outlined below and detailed in the Appendix.

Definitions

Define cooperation as any exchange of data among agents, regardless of whether the exchange is productive or not. If the agents are computer-based, then all the different ways in which they can exchange data can be represented by data-flows [19].

Consider the data flow of a problem-solving organization. Each memory in this data flow is dedicated to some subproblem of the overall problem to be solved, and contains a population of trial-solutions to this subproblem. The dynamics of these populations are determined by their initial values and by the agents that act on them.

Define:

- the effectiveness of a memory to be the double: (5m, vm), where 5m is the highest quality solution that will appear in the memory, and vm is the expected speed of this appearance.
- a memory to be a primary memory if its subproblem is the same as the overall problem.
- a memory to be a cyclic memory if all the agents that write to it also read from it. (Note: any memory in a strongly cyclic data flow is well approximated by a cyclic memory. After all, no memory can see any more of its containing data flow than the agents that read from or write to it. If the data flow is strongly cyclic, then most, if not all, the agents that do one of these things can also be thought of as doing the other. For instance, the entire subgraph that begins with agent-AI and ends with agent Dec in Fig 2 (d), can, from the perspective of the partial tour memory, be replaced by a single super-agent that both reads from, and writes to, the partial tour memory.)
- a memory to be scale-effective if it is cyclic and there are agents that cause 6m to improve monotonically when these agents are added to the memory;
- a data flow to be scale-effective if at least one of its primary memories is so.
- a cyclic memory to be a constant drift memory (CDM) if the dynamics of each of its **trial-**

solutions are described by a Markov chain of the sort shown in Fig. 5.

Some of the more important structural features of a CDM are:

- The members of the initial population of solutions in a CDM are chosen randomly from the space of all possible solutions that can be stored in the memory. The size of this population is small in comparison to the size of the space of all possible solutions.
- A path (a connected sequence of solutions) is developed from each of the members of the initial population by the combined actions of constructors and destroyers. A constructor, when it chooses to act on a path, lengthens it by adding a point to its end; this point is a modification of its immediate predecessor. A destroyer, when it chooses to act on a path, shortens it by removing a point from its end. '
- The paths are developed concurrently. The total time required for the development of each path is the sum of the time that agents spend actually working on the path (adding or erasing points) plus the time by which they are delayed in their work. These delays are of three types: synchronization delays that occur when an agent must pause in order to satisfy a synchronization or precedence constraint in the organization's control structure, communication delays that occur when an agent must wait for the delivery of the data it needs, and resource contention delays that occur when an agent must wait for the computers it requires.
- The CDM and its agents are implemented in a distributed network of computers. Each agent is assigned a computer for its exclusive use, so there are no resource contention delays. Moreover, this computer is sized so that each agent requires the same amount of time for an action as every other agent.

Consider any CDM. Let:

C be the set of construction agents that acts on the CDM.

Q_c be the set of operators contained in C.

D be the set of destruction agents that acts on the memory

S be the space (set) of all possible solutions, good and bad, that can be stored in the memory.

5 be an indicator of solution-quality, such that 6 increases as solution-quality

increases.

G§ be the subset of S that contains all the solutions of quality 8 and better.

N be the size of the initial population of solutions stored in the CDM. Assume that the members of this population are chosen randomly from S and that N is always small in comparison with the size of S.

T_5 be the expected amount of time for the population of solutions to evolve at least one solution of quality 5 or better. G_5 is said to be reachable if T_5 is finite.

5m be the greatest value of 5 such that G_6 is reachable,

vm = $1/T_{5m}$ be the expected speed with which $G\$_m$ is reached.

d(y) be the distance of y from G5, where y is any solution in S_f and d(y) is the minimum number of operations needed to convert y into a member of Gg.

P(S) be a partition of S into regions S_o, S_{1f}...$_f$ S^, such that S_n contains all the solutions that are at a distance of n from G5, as in Fig. 6.

H be the set of all the paths in S.

F(H) be a fuzzy partition of H into regions of desirable and undesirable paths, as in Fig. 7

H_o be the subset of H that the destroyers can recognize and erase.

p be the amount of time required for each agent to take one action.

T_{syn} and T_{com} be the expected synchronization and communication delays experienced by agents in developing a successful path (one that reaches G_{5m}).

Consider s_i and s_{i+1}, the two latest points in any developing path. Let:

p, q and r be the probabilities that s_{i+1} is closer, further and at the same distance, respectively as Sj from G_5.

p_c, q_c and r_c be the values of p,q and r when the destroyers are disabled.

Pd» 3d and rd be the conditional probabilities that Sj^will be destroyed, if it is considered for destruction and if it is further, closer and at the same distance, respectively, as Sj from G5.

\{S\) = p - q be the overall drift of the CDM at s_x\ X^Sj) = p_c - q_c, be the component of drift contributed by the constructors; and X^Sj) = p_d - q_d be the component of drift contrib-

uted by the destroyers. (In a CDM, X is constant, but in other types of memories, X could vary over S.)

A(S) be the space of all the $A_i(S_j)$. Note: in a CDM, X, X_c and X_d are constants for all points at finite and non-zero distances from the goal.

Convergence Conditions And Causal Relations

For any CDM, it can be shown (see the Appendix) that if:

- X is positive, and
- there exists a finite K such that $S_k = 0$ or $S_k \subset H_D$ for all $k > K_f$

then G_5 is reachable.

In other words, any mix of agents that makes the outer regions of S either empty or inaccessible, and makes the drift at all accessible points positive, will produce a solution of quality 5 or better.

It can also be shown that the causal relations among the variables of a **CDM** are as depicted in Fig. 8.

CDM Behavior

The succeeding material applies the convergence conditions and causal relations of CDMs to determine and explain their behaviors along the six dimensions used earlier in describing observed behaviors of A-Teams.

1. Diversity

P(S), the partition of the solution space (Fig. 6), can be reconfigured in two ways: by increasing 6 which causes a migration of points to the outer regions, and by increasing the number of construction operators which causes a migration in the opposite direction. The latter migration is strongest when the new construction operators contain new and powerful skills, in other words, the solution space contracts about a goal set as the range of skills of the construction operators increases. Notice that the addition of certain

destroyers has a similar effect. Of course, destroyers cannot influence the individual distances of solutions from a goal set. But they can make the outer regions of the solution space essentially inaccessible by quickly erasing any solutions that happen to fall there. In effect, they can truncate the solution space, reducing the average distance of the remaining solutions from any goal set, and thereby, causing the accessible part of the solution space to contract about that goal set.

2. Scale

For a CDM to be scale-effective and for solutions of arbitrarily high quality to be reachable, there must be agents that: a) can be added to the CDM, b) make S^ empty or inaccessible, and c) leave the overall drift, *X,* positive for all points that are accessible and not in the goal.

The value of *X* depends on the values of *7*^ and X^, the drifts of the individual constructors and destroyers. These individual drifts are measures of selection acuity. A constructor has a positive value for X_c when the solutions it selects to work on, are moved closer to the goal more often than further away. A destroyer has a positive value of *7*^ when the its decisions to erase solutions are correct more often than wrong.

Expressions for the dependence of *X* on X_c and X^ can be found in the Appendix and visualized with the aid of curves of the sort shown in Rg. 8. This dependence is such that destruction has little effect on overall drift when the construction agents make relatively few selection errors ($X_c > 0.2$). But when construction agents are likely to make numerous selection errors, then destroyers can be used to erase the results of these errors, yielding a high overall drift.

Thus, finding solutions of arbitrarily high quality in a CDM requires neither strategic planning nor coordination. Rather, agents acting independently, without central control, can find these solutions provided only that there is the right mix of agents. One way to achieve such a mix is to: a) include as wide a range of construction knowledge (compiled

into operators) as is available, in order to shrink the outer regions of the solution space, b) design and add destroyers to make whatever remains of the outer regions inaccessible, and c) add destroyers to erase the results of mistaken patterns of construction activity, and thereby make the overall drift positive.

3. Expansion

If agents are non-autonomous, some of their controls are packaged separately in a supervisory system. In such cases, the addition of a new type of agent usually requires a long and painful reengineering of this system. However, when each agent comes with its own complete control system, as is the case with autonomous agents, then no reengineering is necessary. Furthermore, the control system can be customized for the agent's operator: large, complex operators can be paired with appropriately large and complex controls, small operators, with simple controls.

Since the conditions for finding solutions of arbitrarily high quality in a CDM place no restrictions on agent type or granularity large agents can be mixed with small, and general agents with specialists.

4. Duality

In a CDM, construction and destruction are dual processes in the sense that strengths in one can compensate for weaknesses in the other. The evidence is as follows. First, the solution space can be made to contract about any goal set by adding constructors or destroyers. Second, the overall drift is sensitive to both (Fig 8). And third, the subgraphs connecting C and D to 5m and vm in the causal diagram (Fig. 7) are almost symmetric.

5. Population Size

The causal diagram (Fig. 7) indicates that increases in the size of the solution-population do not affect solution-quality but do benefit solution-speed. A more detailed analysis of these benefits (Appendix) shows that they are prone to saturation and affect only the early stages of path-development.

6. Parallelism

In a CDM, the addition of agents improves solution-quality provided these additions leave the overall drift positive. The causal diagram of Fig. 7 indicates that such additions will also improve solution-speed making quality and speed commensurate instead of conflicting attributes, if the completely solid paths from C and D to quality and speed dominate the paths containing broken arcs.

Since synchronization delays are nonexistent for autonomous agents, the exclusive use of such agents causes the T_{syn} node to disappear from the causal diagram along with all the broken paths that go through it. In other words, the exclusive use of autonomous agents makes the conditions for commensuration much easier to meet. (In contrast, the performance of organizations with centralized control systems is often dominated by synchronization delays, making it necessary to pay for increases in solution-quality with decreases in solution-speed.)

CDMs **As Models of More Complex Cyclic Memories**

The observed behaviors of A-Teams are similar to the derived behaviors of CDMs. Are the mechanisms that produce these behaviors also the same? In other words, can CDMs be used to understand the internal workings of A-Team memories? We believe so.

Ofcourse, there are structural differences between A-Team memories and CDMs. Specifically, in an A-Team memory: a) an agent may use several old solutions to produce a new one, b) the overall drift, or rate of diffusion of solutions towards the goal, is unlikely to be constant, and c) the agents may have quite different computational times. We believe these differences are unimportant to solution-quality and can be made unimportant to solution-speed by scheduling strategies such as assigning larger computers to the larger agents. Empirical evidence in support of this belief is emerging in the form of devices, designed on the basis of insights obtained from CDMs, that work in real A-Teams. We feel that it is only a matter of time before theoretical evidence appears in the form of a dilation of the set of CDMs that preserves their behavior but eliminates the structural dif-

ferences with A-Teams.

IV SUMMARY

An A-Team is a sequence of increasingly complex data flows. Each of these data flows consists of sets of memories and agents connected into a strongly cyclic network. The memories represent a non-hierarchic decomposition of the problem to be solved-each memory is dedicated to one subproblem; some of these subproblems may have only a small overlap with the original problem or even be more complex than it.

The later data flows in an A-Team are usually obtained by adding loops to earlier ones, that is, by expanding earlier decompositions and agent-sets. *

In the operation of a data flow, trial-solutions to subproblems are produced by agents and accumulate in memories to form populations. Agents cooperate by working on one another's solutions. This cooperation can take two forms: construction agents add solutions to populations; destruction agents eliminate solutions from populations. In effect, the destroyers can make parts of a solution space inaccessible to constructors by eliminating any solutions that happen to fall in those parts.

Each autonomous agent consists of an operator, a selector and a scheduler. The skills of the agent to create or modify solutions are contained in its operators, the intelligence with which it applies these skills is contained in its selector and scheduler.

The above features in combination-non-hierarchic problem decompositions, populations of solutions, autonomous constructors and destroyers, and strongly cyclic data flows-make A-Teams unusual, perhaps unique, among synthetic problem-solving systems. How can they be understood and made to work?

The observed properties of A-Teams are similar to the derived properties of CDMs (constant drift memories). There is reason to believe that their causal mechanisms are also similar. This being so, the conceptual aids and analytical results for CDMs can be applied

to understand and better design the strongly cyclic data flows that constitute A-Teams.

The most useful of the conceptual aids are P(S), F(H), A(S) and the diagram of Fig.8. P(S) partitions all solutions by their distances from the goal; F(H) partitions all possible patterns of construction activity by their desirability; A(S) is a field of the net rates at which solutions will drift or diffuse towards the goal; and the diagram distinguishes monotonic causal relations from non-monotonic ones. As such, P(S) provides a view of what the set of construction-operators in a data flow can do, if they are controlled perfectly; F(H) provides a view of the mistakes the construction-operators can make, if they are controlled imperfectly; A(S) shows how well the constructors and destroyers will actually do when they are working together; and the causal diagram provides insights into how to improve performance.

The task of reaching the goal (finding a solution of quality 5 or better in a reasonable amount of time), can be broken into two sub-tasks: a) make the outer regions of P(S) empty or inaccessible, and b) make *X(s),* the net drift at point s, positive and as large as possible for all accessible values of s. Any mix of construction and destruction agents that covers these tasks will, if the agents are allowed to work on one another's results, reach the goal. That such mixes exist has been amply demonstrated, at least for optimization. Finding new mixes, however, remains something of a problem. There are, as yet, no automatic procedures, only guidelines. Perhaps the most important of these are:

- use autonomous agents. They have no synchronization delays, and therefore, can provide dramatic speed-ups when allowed to work in parallel. In addition, are easier to add to data flows than non-autonomous agents.
- design agents to encapsulate the relevent knowledge in its naturally occurring chunks. Put knowledge on what to do and where to search into construction agents; knowledge on what to undo and where not to continue to search, into destruction agents; large chunks of knowledge into large agents; small chunks into small agents.
- Include all the best and most powerful construction operators available. The resulting contractions in P(S) invariably exceed those produced by mixes that include only weak construction operators, such as crossover, or a single powerful operator, such as a

state-of-the-art-branch-and-bound.

- use destroyers to increase the net drift by making poor solutions inaccessible and by erasing ineffective patterns of construction over the remaining solutions. In other words, arrange for the destroyers to know or learn about poor solutions and ineffective patterns. This is perhaps the most difficult guideline to implement.

V REFERENCES

[1] G.F. Oster and E.O. Wilson, "Caste and Ecology in the Social Insects," Princeton University Press, Princeton, NJ, 1978.

[2] A. Kerr, Jr., "Subacute Bacterial Endocardites," Charles C. Thomas, Springfield, IL, 1955.

[3] "Handbook of Genetic Algorithms," edited by L Davis, Van Nostrand Reinhold, 1991

[4] H. P. Nii, "Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures, Parts I and II, *At Magazine,* 7:2 and 7:3,1986.

[5] S. Kirkpatrick, C D. Gelatt, and M.P. Cecchi, "Optimization by Simulated Annealing," *Science,* Vol. 220, Number 4598, May, 1983.

[6] F. Glover, Tabu Search-Parts I and II," *ORSA Journal of Computing,* Vol. 1. No. 3, Summer 1989 and Vol. 2, No. 1, Winter 1990.

[7] P.S. de Souza and S.N. Talukdar, "Genetic Algorithms in Asynchronous Teams," *Proceedings of the Fourth International Conference on Genetic Algorithms,* Morgan Kaufmann, Los Altos, CA, 1991.

[8] R.W. Quadrel, "Asynchronous Design Environments: Architecture and Behavior," Ph. D. dissertation, Department of Architecture, Carnegie Mellon University, Pittsburgh, PA, 1991.

[9] S. Murthy, "Synergy in cooperating agents: designing manipulators from task specifications," Ph.D. dissertation, Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 1992.

[10] C.L. Chen, "Bayesian Nets and A-Teams for Power System Fault Diagnosis," Ph. D. dissertation, Electrical and Computer Engineering Department, Carnegie Mellon University, Pittsburgh, PA, 1992.

[11] S. N. Talukdar, V.C. Ramesh, "A parallel global optimization algorithm and its

application to the CCOPF problem/[1] *Proceedings of the Power Industry Computer Applications Conference,* Phoenix, May, 1993.

[12] S. Lin and B.W. Kernighan, "An Effective Heuristic Algorithm for the Traveling-Salesman Problem," *Operations Research,* Vol. 21,1973, pp. 498-516.

[13] M. Held and R.M. Karp, The Traveling-Salesman Problem and Minimum Spanning Trees," *Operations Research,* Vol. 18,1138-1162, 1970.

[14] P. de Souza, "Asynchronous Organizations for Multi-Algorithm Problems," Ph. D. dissertation, Dept. of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 1993.

[15] C. K. Tsen, "Solving Train Scheduling Problems Using A-Teams," Ph.D. dissertation, Electrical and Computer engineering Department, CMU, Pittsburgh, 1995.

[16] S. Y. Chen, S. N. Talukdar, N. M. Sadeh, "Job-Shop-Scheduling by a Team of Asynchronous Agents," *IJCAI-93 Workshop on Knowledge-Based Production, Scheduling and Control,* Chambery, France, 1993.

[17] J. A. Lukin, A. P. Gove, S. N. Talukdar and C. Ho, "An Automated Probabilistic Method for Assigning Backbone Resonances of (13^C, 15^N)-Labelled

[18] S. R. Gorti, S Humair, R. D. Sriram, S. Talukdar, S. Murthy, "Solving Constraint Satisfaction Problems Using A-Teams," to appear in *AI-EDAM.*

[19] S. N. Talukdar and P. S. de Souza "Insects, Fish and Computer-Based Super-Agents," *Systems and Control Theory for Power Systems,* edited by Chow, Kokotovic and Thomas, Vol. 64 of the Institute of Mathematics and its Applications, Springer-Verlag, 1994.

[20] J. H. Kao, J. S. Hemmerie, R P. Prinz, "Asynchronous-Teams Based Collision Avoidance in PAWS," EDRC Report, Carnegie Mellon University, June 1995.

[21] P. Krolak and W. Felts, "A Man-Machine Approach Toward Solving the Traveling Salesman Problem," *Communications of the ACM,* Vol. 14, No. 5, May 1971.

[22] M. Grotschel, "Polyedrische Kombinatorik and Schnittebenverfahren," Preprint No. 38, Universitat Augsburg, 1984.

[23] M. Padberg and G. Rinald, "Optimization of a 532-city Symmetric Traveling Salesman Problem," *Operations Research Letters,* Vol. 6, No. 1, March 1987.

[24] S. N. Talukdar, S. S. Pyo and T Giras, "Asynchronous Procedures for Parallel

Processing," IEEE Trans, on PAS. Vol. PAS-102. NO 11, Nov. 1983.

[25]P. Avila-Abascal and S. N. Talukdar, "Cooperative Algorithms and Abductive Causal Networks for the Automatic Generation of Intelligent Substation Alarm Processors", Proceedings of ISCAS-96

[26] J. Rachlin, F. Wu, S. Murthy, S. Talukdar, M. Sturzenbecker, R. Akkiraju, R. Fuhrer, A. Aggarwal, J. Yeh, R. Henry, R. Jayaraman, "Forest View: A System For Integrated Scheduling In Complex Manufacturing Domains," IBM report, 1996.

[27] H. Lee, S. Murthy, W. Haider, D. Morse, "Primary Production Scheduling at Steel making Industries," IBM Journal of Research and Development, vol. 40, no. 2, pp231-252, 1996

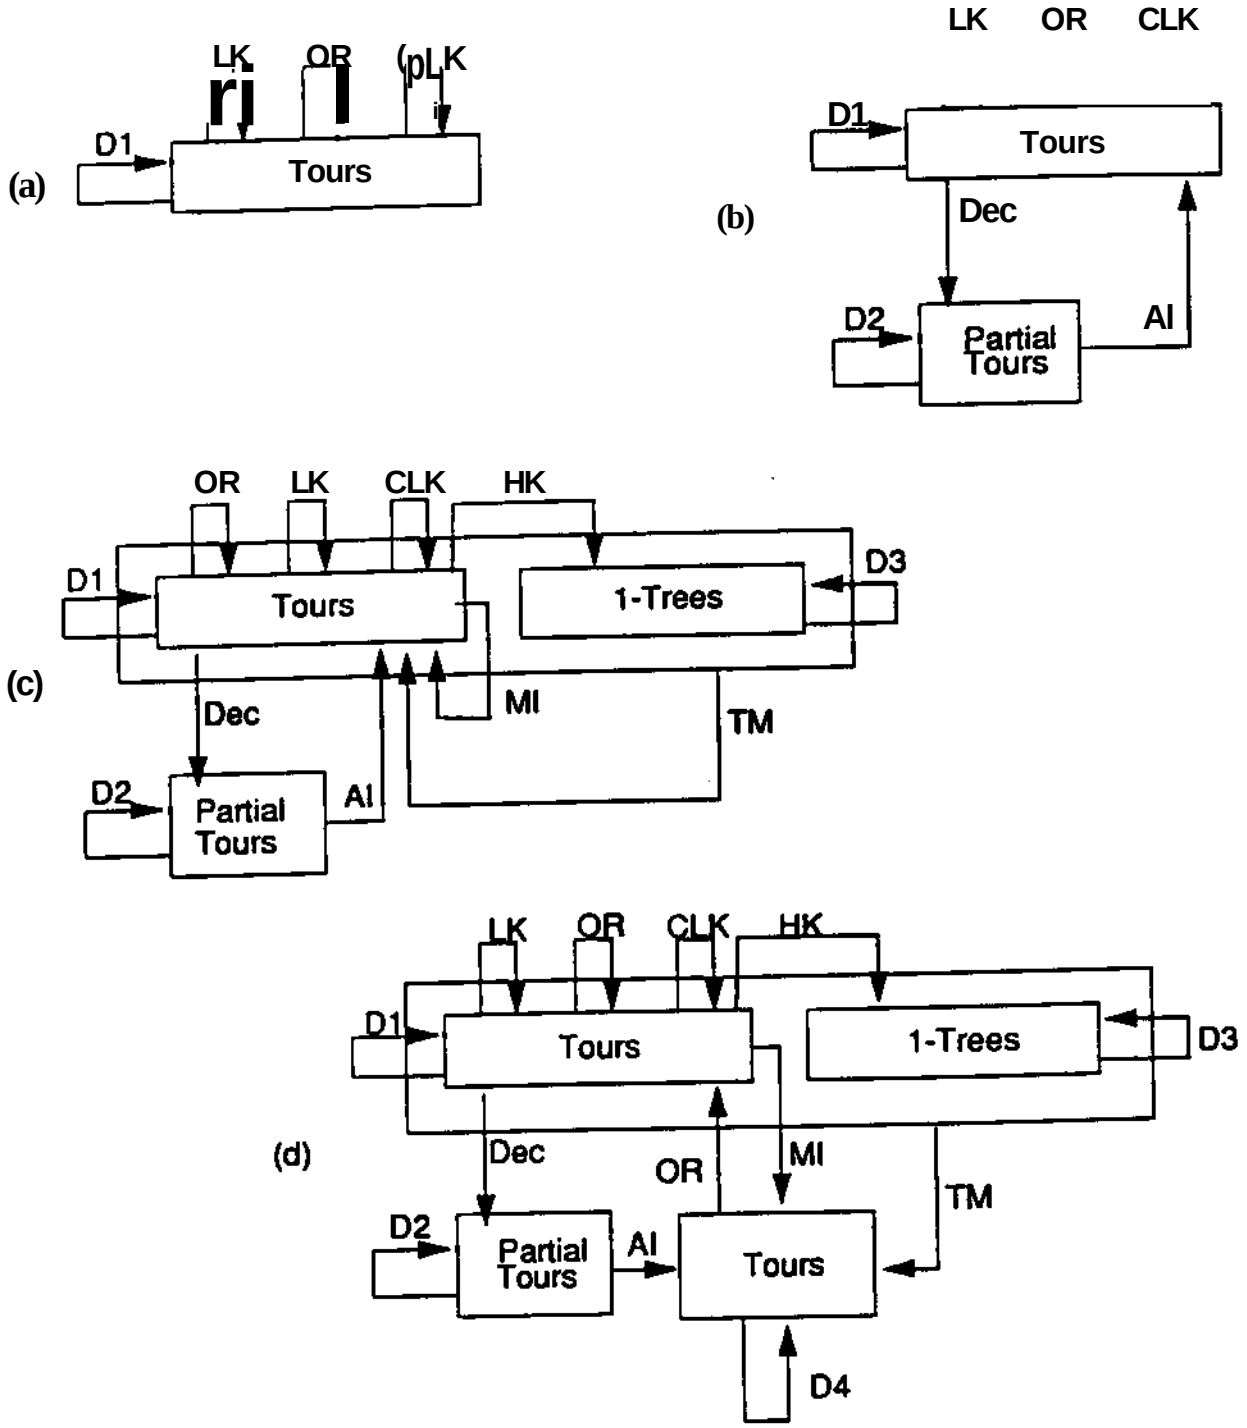

Fig. 1. Four data-flows from an A-Team for the traveling salesman problem. D1-D4 are destruction agents (they eliminate members from solution populations). The other agents are constructive (they add members to solution populations

LK	Lin-Kernighan, one of the longest and most powerful algorithms available [12]
CLK	a shorter and simpler version of LK [14]
OR	Or-Opt_f a moderately complicated algorithm [14]
AI	Arbitrary Insertion, a very short and simple algorithm [14]
HK	Held-Karp, an algorithm for converting tours into 1-Trees [13]
Dec	a deconstructor that produces a partial tour from the common edges of two complete tours [14]
MI	a mixing algorithm that combines two tours to get one [14]
TM	a mixing algorithm that combines a tour with a 1-tree to give a new tour [14].

Fig. 2. A sample of algorithms for the traveling salesman problem.

DATA FLOW (See Fig. 1 for details)	PERFORMANCE A: the difference in length between the best tour that could be found and the optimum tour. T_{5m}: computation time with all the agents sharing one computer (a DEC 5000)							
	Krolak 24 100 cities [21]		LK318 318 cities [14]		PCB 442 442 cities [22]		ATT 532 532 cities [23]	
	A (%)	(sec)	A (%)	T_{8m} (hrs)	A (%)	T_{8m} (hrs)	A (%)	(hrs)
(a)	0	35	1.27	2.9	1.20	4.2	0.87	7.5
(b)	0	39	1.13	2.4	0.89	3	0.47	6.8
(c)	0	39	0.06	1	0.26	4.8	0.40	14
(d)	0	13	0	1.5	0.01	3.5	0.06	13

Fig. 3: Results from applying the data flows of figure 2 to four TSP problems. The results are averages over 15 runs. Each run was terminated when improvements in the tours ceased. All the agents of each A-Team were made to share a single computer- a DEC 5000.

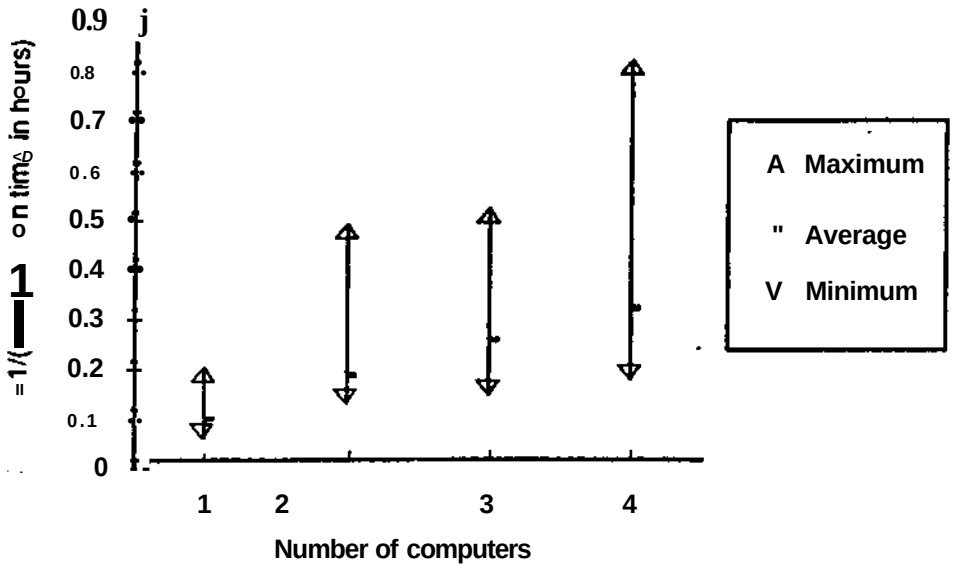

Fig. 4: Plots of speed vs. number of computers for problem ATT532 and the data flow of figure 1 (d). The average, maximum and minimum speeds were for 15 runs.

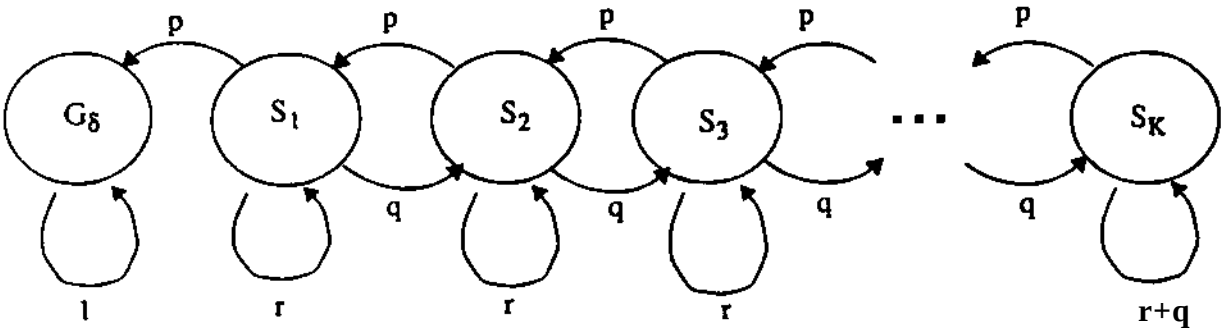

Fig. 5. A Markov chain. Nodes represent solutions states, arcs represent transition probabilities. Consider a trial solution in state S_n. The next agent to work on this solution has a probability p of converting the solution to a solution in S_{n-1}, a probability q of converting it to state S_{n+1} and a probability r of leaving its state unchanged

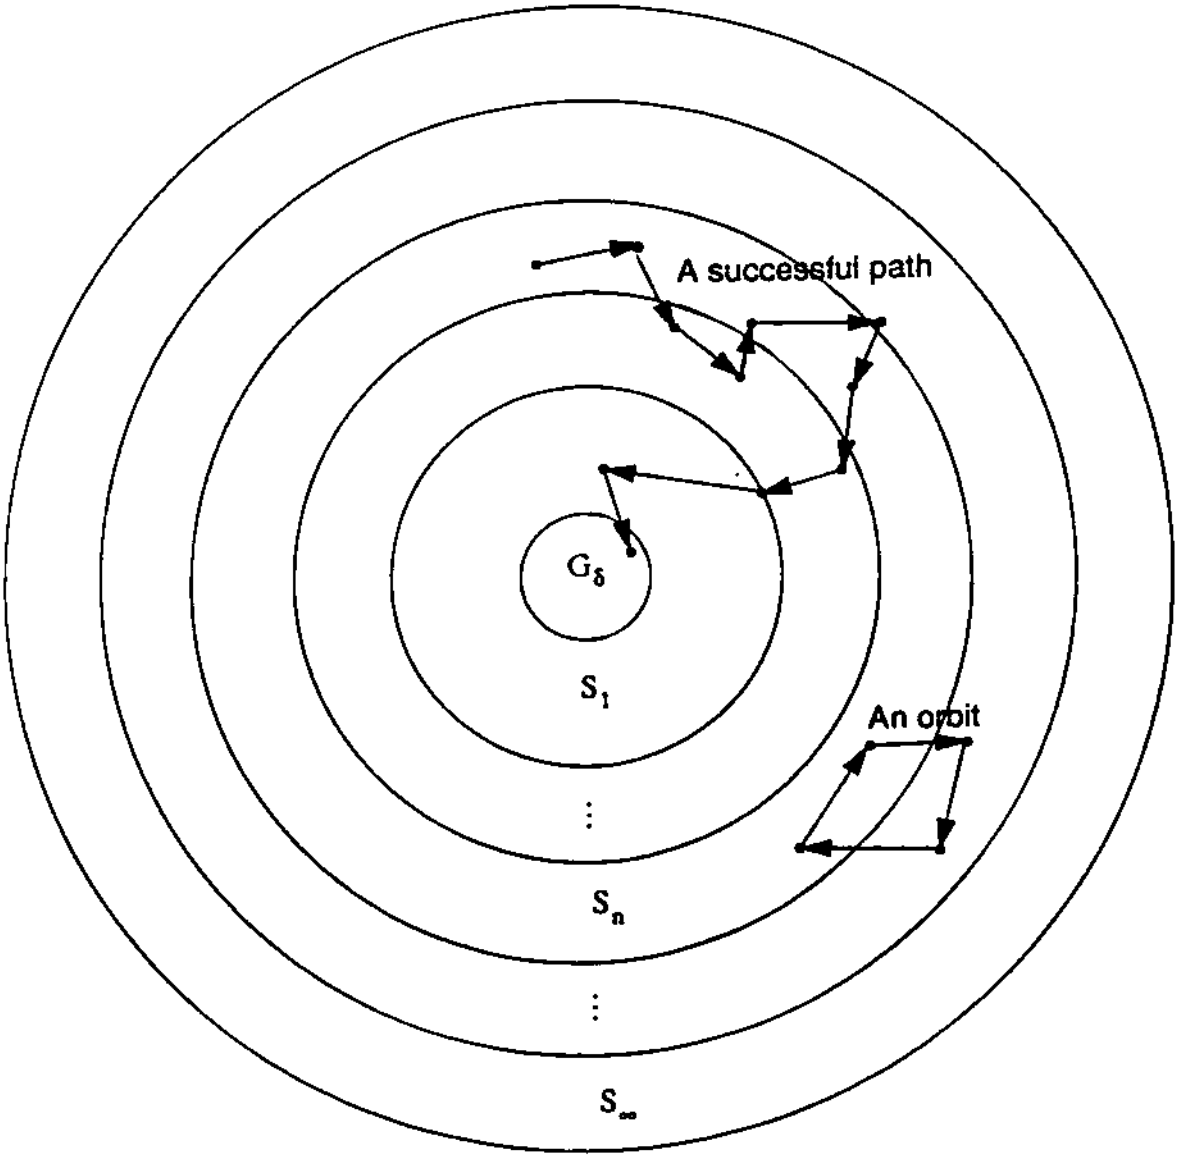

Fig 6. The space of solutions is partitioned into regions so that all the points in S_n are n-constructive operations from the goal space G5. These regions, particularly the outer ones, depend on both G_5 and C. As G5 contracts, the outer regions expand; as C expands, the outer regions contract.

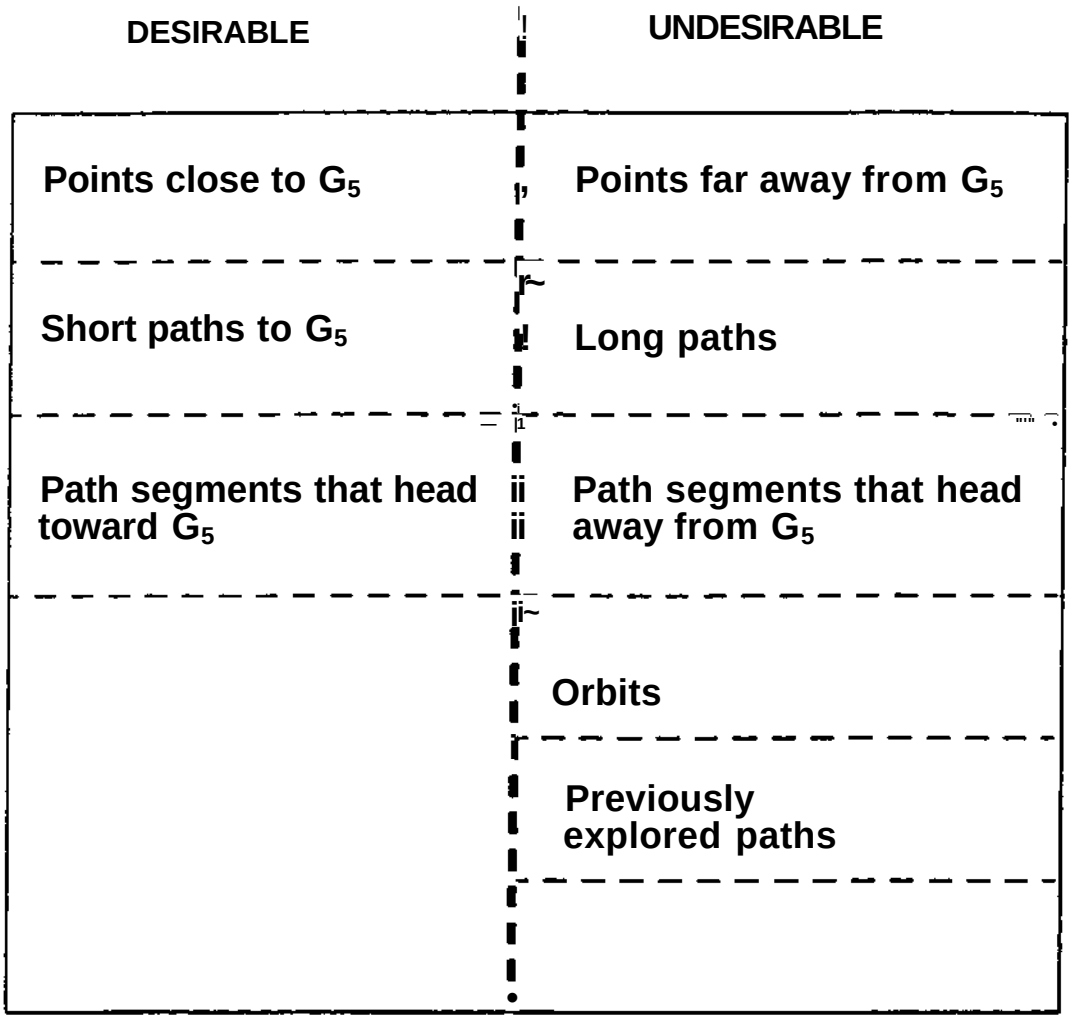

Fig 7: H, the set of all possible **paths** in S, **can,** in **principle, be partitioned into desirable paths and undesirable paths. The function of the destroyers is to recognize and erase** undesirable **paths before** the **constructors have wasted a great deal of time on their** development.

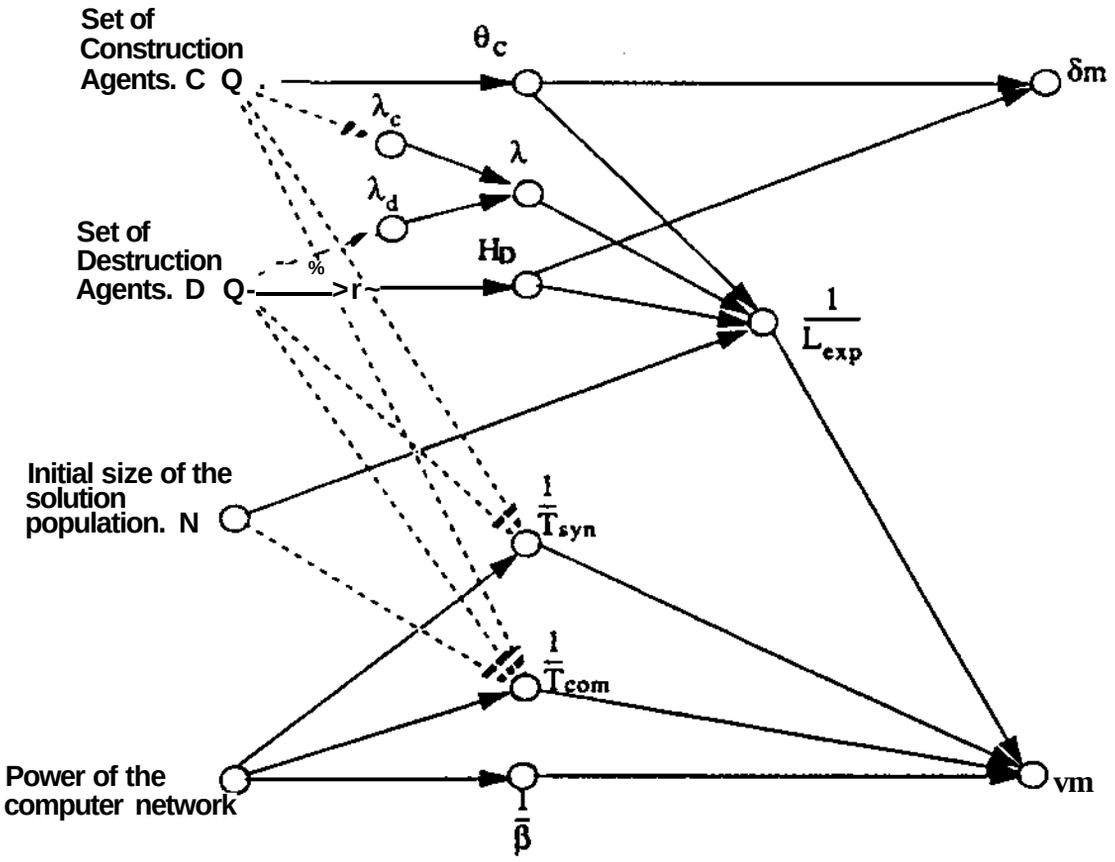

Fig. 8: Causal relations of a CDM. Each solid arc denotes a monotonically-increasing relationship. For instance the solid arc between N and $\frac{1}{L_{exp}}$ means that $\frac{1}{L_{exp}}$ increases monotonically with N.

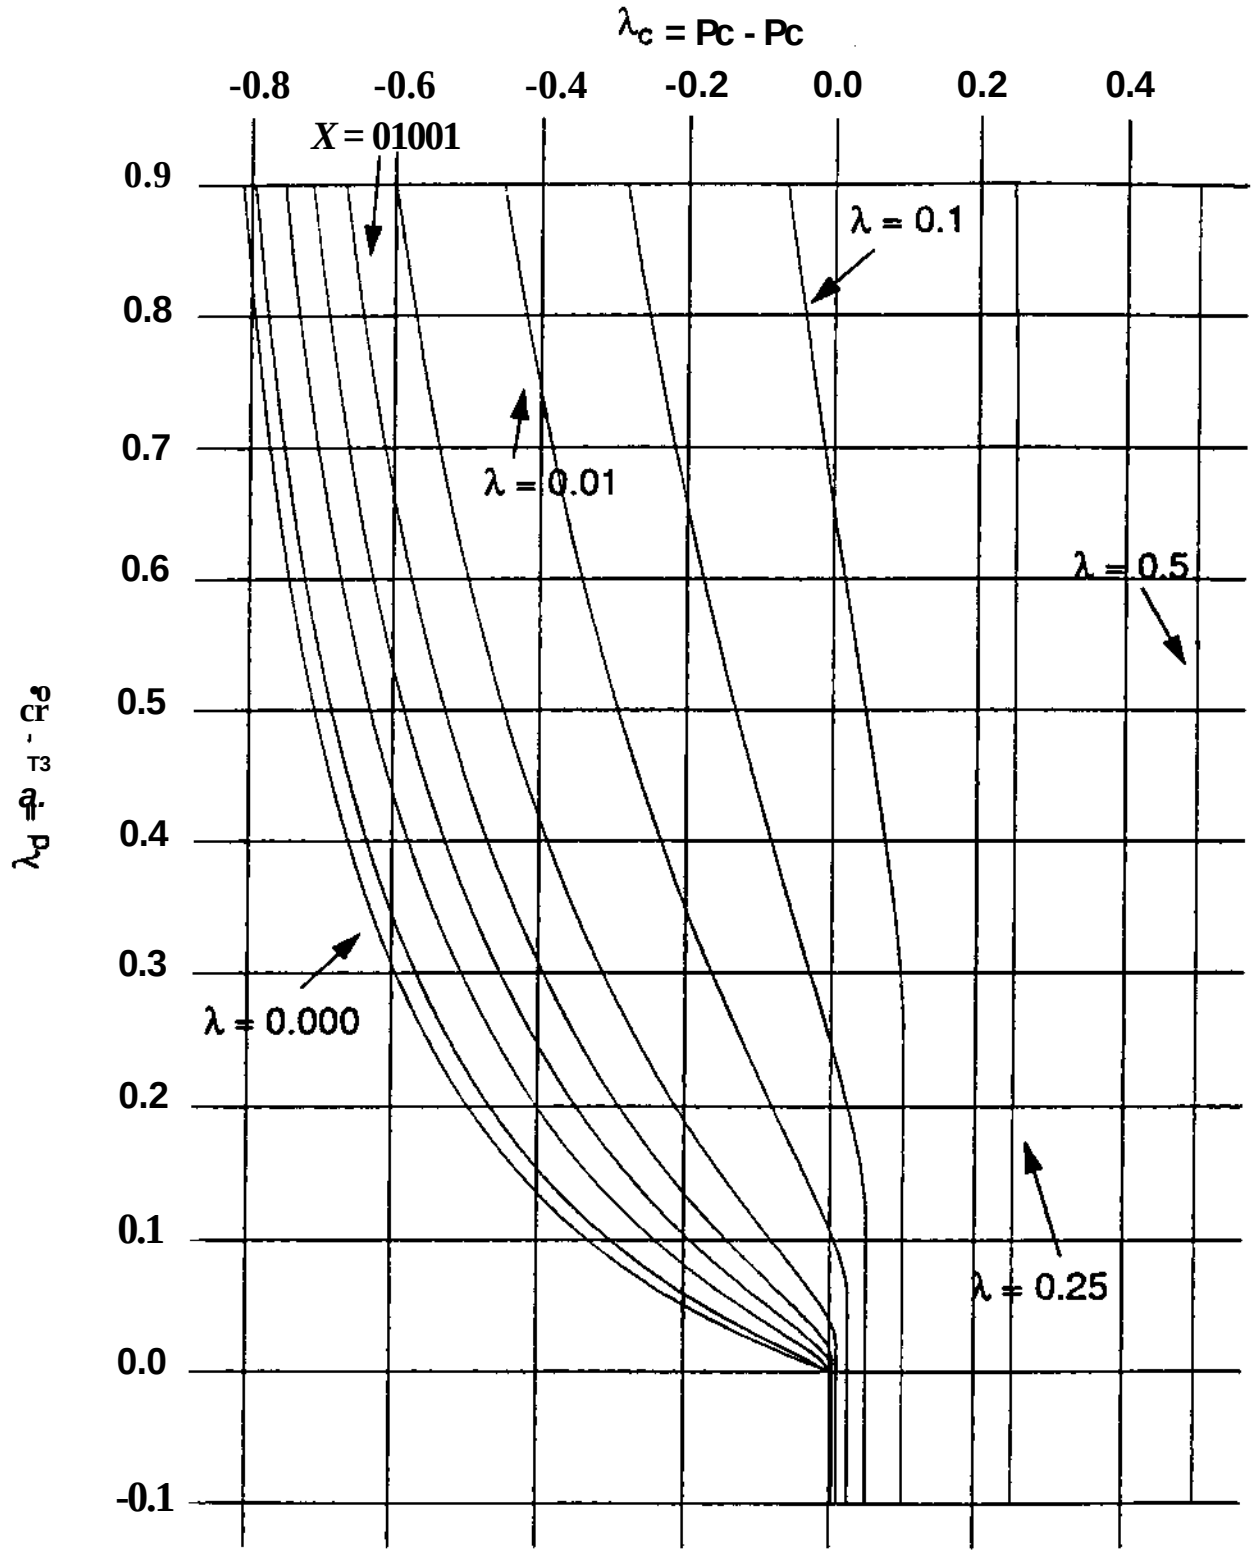

Fig. 9: Iso-drift curves for one set of values of q_j, r_d and r_c. Notice how *X*, the overall drift depends on X_c, the construction drift, and *X$*, the destruction drift

APPENDIX

In this appendix we explore CDMs in greater detail.

Definitions:

Let:

C be the set of construction agents that acts on a CDM

0c be the set of algorithms contained in C.

D be the set of destruction agents that acts on a CDM

S be the space (set) of all possible solutions, good and bad, that can be stored in a CDM.

5 be an indicator of solution-quality such that 8 increases as solution-quality increases.

G_5 be the subset of S that contains all the solutions of quality 5 and better.

N be the size of the initial population of solutions stored in a CDM. (In real problems, N is always small in comparison to the size of S.)

T_s be the expected amount of time for the population of solutions to evolve at least one solution of quality 5 or better. G_5 is said to be reachable if T_8 is finite.

5m be the greatest value of 5 such that G_5 is reachable,

vm = $1/Tg_m$ be the expected speed with which Gg_m is reached.

d(y) be the distance of y from G5, where y is any solution in S, and d(y) is the minimum number of construction-operations needed to convert y into a member of G_5.

S_n be the subset of S containing all the solutions that are at a distance of n from G5, as in Fig. 6.

H be the power set of S (the family of all the subsets of S).

H_D be the subset of H that are recognized and erased y the destroyers in D.

p, q and r be the constant probabilities that the latest edge in any developing path will be a progressive, regressive or neutral edge, respectively; where a progressive edge moves the path's end closer to G5, a regressive edge moves it further away and a neutral edge leaves it at the same distance.

Pc Qc and rc be the values of p,q and r when the.destroyers are disabled.

p_d, q_d and r_d be the conditional probabilities that a regressive edge, if considered for destruction, will be destroyed; that a progressive edge, if considered for destruction, will be destroyed; and that a neutral edge, if considered for destruction, will be destroyed.

X = p-q be the overall drift of the CDM; $X_c = p_c - q_c$, be the drift of the constructors; and $7^{\wedge}$ = Pd " cid be the drift of the destroyers.

P be the amount of time required for each agent to take one action.

T_{syn} and T_{com} be the expected synchronization and communication delays experienced by agents in developing a complete path (one that reaches Gs_m).

Calculating **the overall drift *X* of the system.**

Consider a single path through S as it is developed by the constructors and destroyers. We assume that in each step a constructor which will extend the path is chosen with probability x and a destroyer which may shorten the path is chosen with probability 1 -x. *X* can then be computed from x, p_c, q_c, r_c, p_d, q_d and r_d.

Let:

e denote the edge most recently added to the path.

k_{pro}» k_{reg}» k_{neu} denote the probabilities that e is destroyed some time in the future given that it is a progressive, regressive or neutral edge respectively.

The most recently added edge e can be destroyed in two ways. Either it is destroyed before any other edge is added to the path, or it is destroyed after some other edge has been added. In the latter case the new edge has to be destroyed before e can be considered for destruction again. If we assume the A-team will run for a very large number of iterations, then, once the new edge is destroyed, the system will be in exactly the same state as when e was first added. Hence we get the following recursive formulae for k_{pro}, k_{reg} ^ ^{n}d K_{neu}.

$$k_{pro} = (1-x)\, q_d \sum_{n=0}^{\infty} (1-x)^n (1-q_d)^n$$

$$+ \left[1 - (1-x)\, q_d \sum_{n=0} (1 - (1-q_d)^n \right] (p_c k_{pro} + q_c k_{reg} + r_c k_{ncu})\, k_{pro}$$

$$= \frac{(1-x)\, q_d}{1-(1-x)(1-q_d)} + \left(1 - \frac{}{1-(1-x)(1-q_d)} \right) (p_c k_{pro} + q_c k_{reg} + r_c k_{neu})\, k_{pro}$$

$$k_{reg} = \frac{(1-x)\,p_d}{1-(1-X)\,(1-p_d)} + \left(1 - \frac{(1-x)\,p_d}{1-(1-X)\,(1-p_d)}\right)\left(p_c k_{pf} + q_c k_{reg} + r_c k_{ncu}\right) k_{rcg}$$

$$k = \frac{(1-x)\,r_d}{\ldots} + \left(1 - \frac{(1-x)\,r_d}{\ldots}\right) \ldots$$

The overall drift of the A-team then becomes

$$\lambda = x\,(p_c\,(1-k_{pro}) - q_c\,(1-k_{reg}))$$

Fig. 9 shows X as a function of p_c, q_c, r_c, p_d, q_d and r_d for optimal x.

Reachability of G_5

If G_5 is small or the set of construction agents is weak, not all solutions in S are at a finite distance from G_5. Let S« be this set of points for which there is no path to G_6. Clearly we cannot guarantee reaching G_5 unless the destroyers make S^ inaccessible.

If:

- X is positive, and
- if the outermost regions of S are either empty or made inaccessible by the destroyers,

that is, if there is a finite K such that for k > K, $S_k = 0$ or $S^{\wedge} \subset HQ$

then:

- G_5 is reachable.

proof:

Let j be the expected number of steps required to get one step closer to the goal.

Then

$$j = 1 + p\,(0) + r\,(j) + q\,(2J) \Rightarrow j = \frac{1}{p-q} = \frac{1}{K} \quad (0)$$

so when $X > 0$ we get a finite j. Starting with a solution in S_n the expected number of steps required to get to G_5 is £. let E[n] denote the expected distance to G_5 from the ran-

domly seeded solution. Since all solutions not made inaccessible by the destroyers are at distance at most K we have E[n] < oo. Hence then expected time to goal from the randomly seeded solution is $\frac{E[n]}{A}$ < oo.

Monotonic relationships. (Fig. 7)

1. Solution Speed v_m.

Let:

U be the subset of S that is not in H_D.

R_n be the residue of S_n, that is, the fraction of points in U that are at distances of n or greater from G5. In other words:

If:

• the destruction agents make the portion of S that is outside U completely inaccessible, preventing paths in U from ever leaving it;

• N starting points are randomly chosen from U_f all points in U being equally likely;

• the best (closest to G_5) of these points is identified and a path from it to G5 is developed by the sequential application of construction agents and destroyers;

Then:

$R_n = 0$ if and only if $S_n \cap U = S_{n+1} \cap U = \ldots = 0$ (1)

R_1, R_2, R_3,... decrease monotonically as the variety of constructive skills in increases, that is, as the number of agents in C increases (2)

$$"mi. = \sum_{n=1}^{\infty} R_n^N \quad 0)$$

$$L_{exp} = \frac{1}{\pounds^n min} \quad (4)$$

where n_{min} is the expected distance of the best starting point from G_5, and L_{exp} is the expected length of the path from this point to G5.

Proof:

Result (1) is obvious from the definition of R_n. Result (2) follows directly from the definition of R_n and the fact that S_n decreases as the skills in Q_c increases.

To see (3) note that the probability that the best of N starting points is in distance n from

the goal equals the probability that all points are at least in distance n and not all points are at least in distance n+1 so

$$\%_{,,} = \pounds\bullet(*:->\pounds,) - \sum_{n=1}^{\infty} nR_n^N - \sum_{n=1}^{\infty} (n-1)R_n^N = \sum_{n=1}^{\infty} R_n^N$$

(4) follows immediately from (3) and (0)

Now when N increases each term in the sum (3) decreases so the expected speed

$v_m = \frac{L_{exp}}{n_{min}}$ with which we reach G_5 increases.

Likewise by (2),(4) an increase in *QQ* or *X* also causes the v_m to increase.

2. Solution-quality, b_m.

S^ depends on 8 and *QQ.* Specifically S^ fills as 8 increases, and empties as 0c expands. Suppose that an improvement of solution quality from 8 to 8-A8 causes S« to fill by the amount AS^. Then sufficient conditions for achieving this improvement are: an expansion of 0_C to empty part of AS,*,, and an expansion of H_D to contain the rest of AS^, all while maintaining *X* > 0.

Age Based Restarting.

For small X, the expected path length to the goal Le_{xp} becomes unmanageably large. Faster convergence to the goal can then often be achieved by terminating all paths and restarting with a new random population of solutions. The decision to terminate a path should be made on the basis of information actually available. One piece of information that can readily be used is the number of operations performed since last restart.

Let $n^\circ = (\overset{0}{n}, n^\wedge \ldots J$ be the distribution for the best member in the set of initial solutions, that is $\overset{0}{n}_n$ = probf best initial solution e S_n] , and let

$$P = \begin{bmatrix} 1 & 0 & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\ p & r & q & 0 & \ldots & \ldots & \ldots & \ldots \\ 0 & p & r & q & 0 & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \end{bmatrix}$$

Then $n^N = n^\circ P^N$ is the distribution of distance to goal of the solution after it has been worked on for N iterations. If after N iterations we haven't reached the goal state the conditional distribution of the distance to the goal is given by

$$U_n^{NT} = \Pr\lfloor \text{sol in } S_n(8,5) \text{ after N iterations) sol not in } G_5 \text{ after N iterations} \rfloor = \frac{1}{\left[\prod_{0}^{N} n\right]}$$

and we may say that the solution has age N.

The average remaining number of iterations to get a solution of age N into G_5 is then

$$\sum_{n=i}^{\infty} \frac{n}{p-q} \&_n^N$$ For certain values of p, q, r and n^0, even for some p > q this remaining number of iterations will for some N exceed the expected time to reach the goal from the initial solution. For these values of p, q, r and n° it is beneficial to use age based restarting. Suppose we always reseed the memory after N iterations. Then the expected number of iterations to reach G_5 is

$$^N\left(\text{exp. \# of restarts } J^+\left(^{ex}P\text{- }^{\#\ of}\text{ iterations to reach } GJ_{o|}G_5 \text{ reached in <-N iterations } 1\right)\right)$$

$$= N\left(\frac{1-\Pi_0^N}{N}\right) + \sum_{=1}^{N} t\Pr(X_{t-1} \notin G_\delta \wedge X_t \in G_\delta)$$

$$= N \quad {}^{n_0}\ j\ {}^{+}\ {}^{t}\ t\frac{\Pr(X_{t-1} \in S_1)\,p}{\Pr(X_N \in G_\delta)} \qquad (*)$$

As a numerical example take p = 0.S2, r = 0, q = 0.48 and assume that the best newly created solution will always start in S_{10}. Then the expected number of iterations to get to G_s without age based restarting is 250. With age based restarting at N = 110 we get

n^ = 0.48862 and expected number of iterations to G5 equal to 171.202 < 250.

Note that the restart threshold N that minimizes the expected time to reach G_5 is not necessarily the first N for which the expected number of remaining iterations exceed the expected number of iterations to G5 from the best initial solution. If the variance on the expected remaining iterations is sufficiently high, and it normally is, it is more advantageous to do a few extra iterations before restarting. The best value for N is found by calculating O for a different values of N. It can be shown that (*) is concave in N for fixed p, q, rand n^u.