

AD-A082 058

KANSAS UNIV LAWRENCE DEPT OF ELECTRICAL ENGINEERING F/G 9/2
SCENE ANALYSIS, HOMOMORPHISM, AND CONSISTENT LABELING PROBLEM A--ETC(U)
JAN 80 R M HARALICK, G ELLIOTT AFOSR-77-3307

UNCLASSIFIED

AFOSR-TR-80-0142

NL

[] of []
AD
A082 058

END
DATE
FILMED
4-80
DTIC

AFOSR-TR- 80-0142

LEVEL 2

11

ADA082058

SCENE ANALYSIS, HOMOMORPHISM, AND CONSISTENT
LABELING PROBLEM ALGORITHMS

by

Robert M. Haralick

DTIC
ELECTE
MAR 19 1980

AFOSR-77-3307

January 1980

FILE COPY

Approved for public release;
distribution unlimited.

19 REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER AFOSR TR-80-0142	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) SCENE ANALYSIS, HOMOMORPHISM, AND CONSISTENT LABELING PROBLEM ALGORITHMS.		5. TYPE OF REPORT & PERIOD COVERED Final rept.	
		6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) Robert M. Haralick Gordon Elliott		8. CONTRACT OR GRANT NUMBER(s)	
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Kansas Dept. of Electrical Engineering Lawrence, KA 66045		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61102F 2304 A2	
		11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Office of Scientific Research/NM Bolling AFB, Washington, D. C. 20332	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 1286		12. REPORT DATE Jan 1980	
		13. NUMBER OF PAGES 85	
		15. SECURITY CLASS. (of this report) UNCLASSIFIED	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
18. SUPPLEMENTARY NOTES			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Scene Analysis Region Interpretation Computer Vision Line Labeling Relation Homomorphism Consistent Labeling			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) In this research project the meaning of image matching in the symbolic sense was explored in detail as well as the nature of the algorithms required to do image matching. It was determined that image matching was, in fact, a relational homomorphism problem and that relational homomorphism was a kind of constraint satisfaction problem that is called consistent labeling. Consistent labeling problems are prevalent in image understanding work; for example, line labeling, and region interpretation are consistent labeling problems. Having realized that the hard computational problems were consistent			

20. Abstract cont.

→ labeling problems, our efforts were concentrated on efficient algorithms for solving them. The results indicate that the forward checking algorithm is the fastest algorithm from all algorithm tried for solving the consistent labeling problem. ←

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

I SUMMARY

In this research project the meaning of image matching in the symbolic sense was explored in detail as well as the nature of the algorithms required to do image matching. It was determined that image matching was, in fact, a relational homomorphism problem and that relational homomorphism was a kind of constraint satisfaction problem that is called consistent labeling.

Consistent labeling problems are prevalent in image understanding work: for example, line labeling, and region interpretation are consistent labeling problems. Having realized that the hard computational problems were consistent labeling problems, our efforts were concentrated on efficient algorithms for solving them.

Most of our work generalized the relaxation operators of Ullman and Waltz and showed how what they produced related to either homomorphisms or consistent labeling problems. After we finished these theoretical generalizations, we began empirical studies measuring the improvement of the relaxation algorithms over standard backtracking.

Although due to AFOSR administrative rules, the grant was not continued at VPI after I left Kansas, we nevertheless continued our research and discovered an algorithm that was not of the relaxation type and which gave considerable improvement over them.

The beginning part of this report briefly describes the papers which have been published as a result of this research. The main part of this report describes the results of the research done at VPI.

ALL INFORMATION CONTAINED HEREIN IS UNCLASSIFIED
DATE 11-13-81 BY 1043
EXCEPT WHERE SHOWN OTHERWISE (75).
DISTRIBUTION STATEMENT A
A. D. HODGSON
Technical Information Officer

II Papers Published

One important process in many artificial intelligence problems is that of making empirical knowledge (coming from data) about the world consistent with theoretical or model knowledge we have about the world. Constraint relaxation techniques have been used with some success in this regard. In a paper entitled "Scene Analysis, Arrangements, and Homomorphisms" presented at the 1977 Workshop on Scene Analysis at Amherst, Massachusetts, and subsequently published as a chapter in the book Machine Vision (Hanson and Riseman, Ed., 1978), we illustrated how the constraint relaxation as applied to scene analysis problems is really a technique for finding homomorphisms from one arrangement (a set of labeled relations) to another. We showed how a few common scene analysis problems are really problems in establishing the relationship between one arrangement and another.

The idea of homomorphism between arrangements and a discrete relaxation algorithm for finding them was discussed in detail in the paper "Arrangements, Homomorphisms, and Discrete Relaxation" (with J. Karten) that appeared in the August 1978 issue of the IEEE Transactions on Systems, Man, and Cybernetics.

In a paper entitled "Structural Pattern Recognition, Arrangements, and Theory of Covers" presented at the 1977 Pattern Recognition and Image Processing Conference, we illustrated how the concept of arrangements can be used in structural pattern recognition as an alternative to syntactic techniques. This paper was substantially expanded, illustrating how syntactic techniques depend on homomorphisms, and published in the June 1978 issue of Pattern Recognition. The paper was entitled "Structural Pattern Recognition, Homomorphism, and Arrangements".

Further generalizing this basically combinatorial problem and the operations to make the tree search more efficient, we began work on the consistent labeling problem. A talk was presented at the 1978 Princeton Workshop on Artificial Intelligence and Pattern Recognition and Image Processing. The entire scope of this work was prepared with Shapiro as a two part paper entitled "The Consistant Labeling Problem" part one of which appeared in the March 1979 issue of the IEEE Transactions on Pattern Analysis and Machine Intelligence and part two of which will appear in the May 1980 issue.

in the appendix of this report

List of Publications

- "Scene Analysis, Arrangements, and Homomorphisms," Workshop on Computer Vision Systems, University of Massachusetts, June 1-3, 1977. (Invited)
- "Structural Pattern Recognition, Arrangements, and Theory of Covers," IEEE Conference on Pattern Recognition and Image Processing, Rensselaer Polytechnic Institute, Troy, New York, June 6-8, 1977.
- "Increasing Tree Search Efficiency for Constraint Satisfaction Problems" (with G. Elliott) International Joint Conference on Artificial Intelligence, Tokyo, Japan, August 1979.
- "The Characterization of Binary Relation Homomorphisms," International Journal General Systems, Vol. 4, 1978, pp. 113-121.
- "Reduction Operations for Constraint Satisfaction," (with Larry S. Davis, Azriel Rosenfeld, and David L. Milgram), Information Sciences, Vol. 14, No. 3, April 1978.
- "Scene Analysis, Arrangements, and Homomorphisms," Chapter contributed to Machine Vision, (Hanson and Riseman, Editors), Academic Press, New York, 1978.
- "Structural Pattern Recognition, Homomorphisms, and Arrangements," Pattern Recognition, Vol. 10, No. 3, June 1978.
- "Arrangements, Homomorphisms and Discrete Relaxation," (with J. Kartus), IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-8, No. 8, August 1978, pp. 600-612.
- "The Consistent Labeling Problem I," (with L. Shapiro), IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 1, No. 2, March 1979, pp.

INCREASING TREE SEARCH EFFICIENCY
FOR
CONSTRAINT SATISFACTION PROBLEMS

by

Robert M. Haralick
and
Gordon Elliott

Department of Electrical Engineering
Virginia Polytechnic Institute and State University
Blacksburg, Virginia 24061

TABLE OF CONTENTS

ABSTRACT	1
I. INTRODUCTION	1
II. SOME PROCEDURES FOR TREE SEARCH	7
II.1 LOOKING AHEAD	7
II.2 PARTIAL LOOKING AHEAD	9
II.3 FORWARD CHECKING	10
II.4 BACKCHECKING	13
II.5 BACKMARKING	13
III. EXPERIMENTAL RESULTS	21
IV. STATISTICAL MODEL FOR CONSTRAINT SATISFACTION SEARCHES .	47
V. THE FAIL FIRST PRINCIPLE	63
V.1 OPTIMIZING THE CONSISTENCY CHECK ORDER IN TREE SEARCHING.	63
V.2 OPTIMIZING TREE SEARCH ORDER	68
VI. CONCLUSION	75
APPENDIX A. COMPUTER PROGRAM LISTINGS	78

ABSTRACT

In this paper we explore the number of tree search operations required to solve binary constraint satisfaction problems. We show analytically and experimentally that the two principles of first trying the places most likely to fail and remembering what has been done to avoid repeating the same mistake twice improve the standard backtracking search. We experimentally show that a lookahead procedure called forward checking (to anticipate the future) which employs the most likely to fail principle performs better than standard backtracking, Ullman's Waltz's, Mackworth's, and Haralick's discrete relaxation in all cases tested, and better than Gaschnig's backmarking in the larger problems.

I. INTRODUCTION

Associated with search procedures are heuristics. In this paper we provide a theory which explains why two heuristics used in constraint satisfaction searches work. The heuristics we discuss can be given a variety of one line descriptions such as:

Lookahead and anticipate the future in order to succeed in the present.

To succeed, try first where you are most likely to fail.

Remember what you have done to avoid repeating the same mistake.

Lookahead to the future in order not to worry about the past.

We will attempt to show that for a suitably defined random constraint satisfaction problem, the average number of tree search operations which employs these principles will be smaller than that required by the standard backtracking tree search.

To begin our discussion, we need a precise description of the constraint satisfaction problem we are attempting to solve by a search procedure. We assume that there are N units (some authors call these variables instead of units). Each unit has a set of M possible values or labels. The constraint satisfaction problem we consider is to determine all possible assignments f of labels to units such that for every pair of units, the corresponding label assignments satisfy the constraints. More formally, if U is the set of units and L is the set of labels, then the binary constraint R can be represented as a binary relation on $U \times L$: $R \subseteq (U \times L) \times (U \times L)$. If a pair of unit-

labels $(u_1, \ell_1, u_2, \ell_2) \in R$, then labels ℓ_1 and ℓ_2 are said to be consistent or compatible for units u_1 and u_2 . A labelling f of all the units satisfies the constraints if for every pair u_1, u_2 of units $(u_1, f(u_1), u_2, f(u_2)) \in R$. Haralick et al. (1978) call such a labeling a consistent labeling.

The problem of determining consistent labelings is a general form of many problems related to artificial intelligence. For example, scene labeling and matching (Barrow and Tenebaum, 1976, and Rosenfeld et al., 1976), line interpretation (Waltz, 1972), edge labeling (Haralick, 1978), graph homomorphisms and isomorphisms (Ullman, 1969), graph coloring (Harary, 1969), boolean satisfiability (Haralick et al., 1978), and proposition theorem proving (Kowalski, 1975) are all special cases of the general consistent labeling problem. Ullman (1966), Waltz (1972), Rosenfeld et al. (1978 and 1979), Gaschnig (1977, 1978 and 1979), and McGregor (1979) attempt to find efficient methods to solve the consistent labeling problem. Knuth (1975) also analyzes the backtracking tree search, which is the basis of most methods used to solve the consistent labeling problem.

For the purpose of illustrating the search required to solve this problem, we choose the N-queens problem, how to place N-queen on an $N \times N$ checkerboard so that no queen can take another. Here, the unit set corresponds to the row coordinates on a checkerboard and we denote them by positive integers. The label set corresponds to the column coordinates on a checkerboard and we denote them by alphabetic

characters. Hence, the unit-label pair (1,A,2,D) satisfies the constraint R, $[(1,A,2,D) \in R]$, since a queen on row 1 column A cannot take a queen on row 2 column D. But, the unit label pair (1,A,3,C) does not satisfy the constraint R because queens can take each other diagonally (see figure 1).

Using the number letter convention for unit-label pairs, Figure 2 illustrates a portion of a backtracking tree trace for the 6-queens problem. Notice how the unit 5 labels A, C, E, and F occur twice in the trace, each time being tested and failing for the same reason: incompatibility with units 1 or 2. These redundant tests can be eliminated if the fact they failed can be remembered or if units 1 or 2 could lookahead and prevent 5 from taking the labels A, C, E, or F. The remembering done by Gaschnig's backmarking (1977) and the forward checking approach described in this paper help eliminate these problems. Notice that once unit 3 takes label E (Figure 1a) the only labels left for units 4 and 6 are incompatible. The forward checking algorithm will not discover this future incompatibility. However, the first time label B is associated with unit 4, there is absolutely no label possible for unit 6. Hence, the search through the labels for 5 and 6 are entirely superfluous and forward ckecking will discover this (Figure 1b). The lookahead procedures (discrete relaxation) of Ullman (1966), Waltz (1972), Rosenfeld (1976), Mackworth (1977), and Montanari (1974) help alleviate the problem illustrated in Figure 1a as well as in Figure 1b.

Section II gives a description of the full and partial looking ahead, forward checking, backchecking, and backmarking procedures. In section III we compare the complexity of these algorithms as they solve the N-queens problem and problems generated randomly. We measure complexity in terms of number of table lookups and number of consistency checks. These results show that standard backtracking is least efficient in most cases and bit parallel forward checking is most efficient for the cases tried.

In section IV, we give a statistical analysis of constraint satisfaction searches and demonstrate the statistical reason why forward checking requires fewer expected consistency checks than standard backtracking. In section V we explore other applications of the fail first or prune early tree search strategies and show that such particular strategies as choosing the next unit to be that unit having fewest labels left and testing first against units whose labels are least likely to succeed reduce the expected number of consistency tests required to do the tree search. Finally, by changing the unit search order dynamically in every tree branch so that the next unit is always the one with fewest labels left, we show experimentally that performance improves for each procedure and that forward checking even increases its computational advantage over the other algorithms.





Figure 1a illustrates how the labeling A,C,E for units 1,2,3 implies that the only labels for units 4 and 6 are incompatible in the 6 queens problem.

Figure 1b illustrates how the labeling A,C,E,B for units 1,2,3,4 implies that there is no label for unit 6 in the 6 queens problem.

```
1  A
   2  A,B
   2  C
     3  A,B,C,D
     3  E
       4  A
       4  B
         5  A,B,C
         5  D
           6  A,B,C,D,E,F
           5  E,F
         4  C,D,E,F
     3  F
       4  A
       4  B
         5  A,B,C,D,E,F
         4  C,D,E,F
```

Figure 2 illustrates a segment of a tree trace that the standard backtracking algorithm produces for a 6 queens problem. No solutions are found in this segment. The entry 2 A,B, for example, indicates that labels A and B were unsuccessful at level 2, but 2 C succeeds when checked with past units, and the tree search continues with the next level.

II. SOME PROCEDURES FOR TREE SEARCH REDUCING

In this section we give brief descriptions of five procedures, and a variation of data structure in one, which can be used within the standard backtracking framework to reduce tree search operations. They are called full and partial looking ahead, forward checking, backchecking, and backmarking. Each of these procedures invests resources in additional consistency tests or data structures at each point in the tree search in order to save (hopefully) more consistency tests at some point later in the tree search.

For ease in explaining these procedures, we call those units already having labels assigned to them the past units. We call the unit currently being assigned a label the current unit and we call units not yet assigned labels the future units. We assume the existence of a unit-label table which at each level in the tree search indicates which labels are still possible for which units. Past units will of course have only one label associated with each of them. Future units will have more than one. The tree search reducing procedures invest early to gain later. Hence, the result of applying any of them in the tree search will be to decrease the number of possible labels for any future unit or reduce the number of tests against past units.

II.1 Looking Ahead

Waltz filtering (Waltz, 1972), a procedure by Ullman (1966), discrete relaxation (Rosenfeld, Hummel, Zucker, 1976), and the

Ψ operator of Haralick et al. (1978) are all examples of algorithms that look ahead to make sure that (1) each future unit has at least one label which is compatible with the labels currently held by the past and present units and (2) each future unit has at least one label which is compatible with one of the possible labels for each other future unit. Looking ahead prevents the tree search from repeatedly going forward and then backtracking between units u and v , $v < u$, only to ultimately discover that the labels held by units 1 through v cause incompatibility of all labels between some unit w , $w > u$, and some past, current, or future unit.

Because looking ahead in this manner cannot remember and save most of the results of tests performed in the lookahead of future units with future units for use in future lookaheads, the full savings of looking ahead are not realized for many problems. A partial look ahead that does not do all the checks of full look ahead will perform better, and one that checks only future with present units (neglects future with futures) will do much better because all tests it performs can be usefully remembered.

The procedure L_A_TREE SEARCH and its associated subroutines CHECK-FORWARD and LOOK_FUTURE (Figure 3 a,b, and c) is a formal description of the full looking ahead algorithm, which can easily be translated into any structured recursive language. U is an integer representing the unit, and will increment at each level of the tree search. It takes on the value 1 at the initial call. F is a one dimension array indexed by unit, where entry $F(u)$ for unit u is the label assigned to u . T and NEW_T are tables, which can be thought

of as an array of lists, $T(u)$ is a list of labels which have not yet been determined to be not possible for unit u . (We implemented T as a 2 dimension array, with the number of entries in each list (or row) stored in the first position of the row. This implementation uses approximately $(\text{NUMBER_OF_UNITS})^2 \times (\text{NUMBER_OF_LABELS})$ words of memory for table storage since there can be NUMBER_OF_UNITS levels of recursion.) The tree search is initially called with T containing all labels for each unit. All other variables can be integers. EMPTY_TABLE and NUMBER_OF_UNITS and NUMBER_OF_LABELS have obvious meanings.

The function $\text{RELATION}(u_1, l_1, u_2, l_2)$ returns *TRUE* if $(u_1, l_1, u_2, l_2) \in R$, otherwise it returns *FALSE*. CHECK_FORWARD checks that each future unit label pair is consistent with the present label $F(u)$ for unit u as it copies the table T into the next level table NEW_T . LOOK_FUTURE then checks that each future unit label pair in NEW_T is consistent with at least one label for every other unit, and deletes those that are not.

In this implementation CHECK_FORWARD and LOOK_FUTURE return a flag, EMPTY_ROW_FLAG , if a unit is found with no possible consistent labels. Thus the next level of the tree search will not be called, otherwise each entry in NEW_T is consistent with $u, F(u)$, and therefore, all the past unit-label pairs.

II.2 Partial Looking Ahead

Partial looking ahead is a variation of looking ahead which does approximately half of the consistency checks that full looking ahead does while checking future with future units. Each future unit-label

pair is checked only with units in its own future, rather than all other future units. Thus partial looking ahead is less powerful than full looking ahead in the sense that it will not delete as many unit-label pairs from the lists of potential future labels. We will, however, see that partial looking ahead does fewer total consistency checks than full looking ahead in all cases tested.

The checks of future with future units do not discover inconsistencies often enough to justify the large number of tests required, and these results cannot be usefully remembered. Since partial looking ahead does fewer of these less useful tests, it is more efficient. A look ahead that checks only future with current or past units can have better performance since these more powerful tests can also be usefully remembered.

The formal algorithm for partial looking ahead is `L_A_TREE_SEARCH` (figure 3a), with the call to `LOOK_FUTURE` on line 5 replaced with an identical call to `PARTIAL_LOOK_FUTURE` (figure 3d).

II.3 Forward Checking

Forward checking is a partial lookahead of future units with past and present units, in which all consistency checks can be remembered for a while. This method is similar to looking ahead, except that future units are not checked with future units, and the checks of future units with past units are remembered from checks done at past levels in the tree search. Forward checking begins with a state of affairs in which there is no future unit having any of its labels inconsistent with any past unit-label pairs. This is certainly true

at the base of the tree search, since there are no past units with which to be inconsistent. Because of this state of affairs, to get the next label for the current unit, forward checking just selects the next label from the unit label table for the current unit. That label is guaranteed to be consistent with all past unit label-pairs. Forward checking tries to make a failure occur as soon as possible in the tree search by determining if there is any future unit having no label which is consistent with the current unit-label pair. If each future unit has consistent labels, it remembers by copying all consistent future unit-label pairs to the next level's unit label table. If every future unit has some label in the unit label table which is consistent with the current unit-label pair, then the tree search can move forward to the next unit with a state of affairs similar to how it started. If there is some future unit having no label in the unit label table which is consistent with the current unit-label pair, then the tree search remains at the current level with the current unit and continues by selecting the next label from the table. If there is no label then it backtracks to the previous unit and the previous label table.

The formal algorithm for forward checking is the Procedure L_A_TREE_SEARCH (Figure 3a) with line 5, the call to LOOK_FUTURE, removed. Forward checking is just looking ahead, omitting the future with future checks.

An improvement in efficiency can be gained in the lookahead type of algorithms by using a data structure for the unit-label tables that is suggested by McGregor (1979). McGregor simultaneously

developed a weaker form of the forward checking algorithm and compared them on subgraph isomorphism problems. In `L_A_TREESEARCH` and `CHECK_FORWARD` the lists `T(u)` or `NEW_T(u)` can be stored as bit vectors, with one bit in the machine word for each possible label (this is essentially a set representation). Lines 4, 5 and 6 in `CHECK_FORWARD` (figure 3b) can then be replaced with the following statement:

`NEW_T(U2) = AND( T(U2), RELATION_BIT_VECTOR(U,L,U2) )`

This single statement replaces a loop, taking advantage of the parallel bit handling capabilities of most computers. `RELATION_BIT_VECTOR (u,l,u2)` returns a bit vector with a bit on in each position corresponding to a label l_2 for which `RELATION (u,l,u2,l2)` would have been true. This is essentially the set $\{l_2 \mid (u,l,u_2,l_2) \in R\}$. Thus `NEW_T (u2)` becomes the set $\{l_2 \in T(u_2) \mid (u,l,u_2,l_2) \in R\}$, precisely what lines 4, 5, and 6 do.

If the number of labels is less than the word length of the computer used, then the relation can be directly stored in an array of size $(\text{NUMBER_OF_UNITS})^2 \times (\text{NUMBER_OF_LABELS})$, and the tables `T` and `NEW_T` will take approximately $(\text{NUMBER_OF_UNITS})^2$ words of storage. Reduced forms of the relation exist for some problems, such as the N-queens problem or the subgraph isomorphism problem, in which only two dimensional tables need be stored and a quick calculation will generate the needed list. The same technique can be applied to the full and partial looking ahead algorithms, but they will not be compared here since the three algorithms will have approximately the same relationships of efficiency, to each other, with or without the improved data structure.

II.4 Backchecking

Backchecking is similar to forward checking in the way it remembers unit label pairs which are known to be inconsistent with the current or any previous unit label. However, it keeps track of them by testing the current unit label only with past unit label pairs and not future ones. So if, for instance, labels A, B, and C for unit 5 were tested and found incompatible with label B for unit 2, then the next time unit 5 must choose a label, it should never have A, B, or C as label possibilities as long as unit 2 still has the label B.

Each test that backchecking performs while looking back from the current unit u to some past unit v , forward checking will have performed at the time unit v was the current unit. Of course, at that time, forward checking will also have checked all future units beyond unit u . Hence, backchecking performs fewer consistency tests, an advantage. But backchecking pays the price of having more backtracking and at least as large a tree as forward checking. Backchecking by itself is not as good as forward checking.

II.5 Backmarking

Backmarking (defined in Gaschnig, 1977, and also discussed in Gaschnig, 1978) is backchecking with an added feature. Backchecking eliminates performing some consistency checks that were previously done, had not succeeded, and if done again would again not succeed. Backmarking also eliminates performing some consistency checks that

were previously done, had succeeded, and if done again would again succeed. To understand how backmarking works, recall that the tree search by its very nature goes forward, then backtracks, and goes forward again. We focus our attention on the current unit u . We let v be the lowest ordered unit to which we have backtracked (has changed its label) since the last visit to the current unit u . Backmarking remembers v . If $v = u$, then backmarking proceeds as backchecking. If $v < u$, then since all the labels for unit u had been tested in the last visit to unit u , any label now needing testing, needs only to be tested against the labels for units v to $u-1$, which are the ones whose labels have changed since the last visit to unit u . That is, the tests done previously against the labels for units 1 through $v-1$ were successful and if done again would again be successful because labels for units 1 through $v-1$ have not changed and the only labels permitted for the current unit u are those which have passed the earlier tests (see figure 4).

The formal algorithm for Backmarking appears in figure 5. It is essentially Gaschnig's algorithm (in Gaschnig, 1977), but modified to find all solutions. The variable U and array F are the same as in looking ahead. $LOWUNIT$ is a one dimensional array of $NUMBER_OF_UNITS$ entries, and $LOWUNIT(i)$ will indicate the lowest level at which a change of label has occurred since the last time the $MARK$ array is set. $MARK$ is dimensioned $NUMBER_OF_UNITS$ by $NUMBER_OF_LABELS$, and $MARK(u, \ell)$ will indicate the lowest level at which a consistency test failed when the unit label pair (u, ℓ) at the current level was last tested against the previous unit

label pairs on previous levels. At any point if $\text{MARK}(u, \ell)$ is less than $\text{LOWUNIT}(u)$ then the algorithm knows that (u, ℓ) has already been tested against the unit-label pairs at levels below the value in $\text{LOWUNIT}(u)$ and will fail at level $\text{MARK}(u, \ell)$, so there is no need to repeat the tests. If $\text{MARK}(u, \ell)$ is greater or equal to $\text{LOWUNIT}(u)$ then all tests will succeed below the level $\text{LOWUNIT}(u)$ and only tests against units at $\text{LOWUNIT}(u)$ to the current unit need be tested.

Before the initial call to BACKMARK, all entries in LOWUNIT and MARK are initialized to 1, and BACKMARK is called with the initial $u = 1$. Since the same MARK and LOWUNIT arrays are used at all levels of recursion of the tree search, approximately $(\text{NUMBER_OF_UNITS}) \times (\text{NUMBER_OF_LABELS})$ words of table storage are needed.

```
1.  RECURSIVE PROCEDURE L_A_TREE_SEARCH(U,F,T);
2.  FOR F(U) = each element of T(U) BEGIN
3.    IF U < NUMBER_OF_UNITS THEN BEGIN
4.      NEW_T = CHECK_FORWARD(U,F(U),T);
5.      CALL LOOK_FUTURE(U,NEW_T);
6.      IF NEW_T is not EMPTY_ROW_FLAG THEN
7.        CALL L_A_TREE_SEARCH(U+1,F,NEW_T);
8.    END;
9.  ELSE
10.    Output the labeling F;
11.  END;
12. END L_A_TREE_SEARCH;
```

Figure 3a

```
1.  PROCEDURE CHECK_FORWARD(U,L,T);
2.  NEW_T = empty table;
3.  FOR U2 = U+1 TO NUMBER_OF_UNITS BEGIN
4.    FOR L2 = each element of T(U2)
5.      IF RELATION(U,L,U2,L2) THEN
6.        Enter L2 into the list NEW_T(U2);
7.      IF NEW_T(U2) is empty THEN
8.        RETURN (EMPTY_ROW_FLAG); /* No consistent labels */
9.    END;
10. RETURN (NEW_T);
11. END CHECK_FORWARD;
```

Figure 3b

Figures 3a, b, and c express the full looking ahead algorithm. Replace the call to LOOK_FUTURE at line 5 in figure 3a with an identical call to PARTIAL_LOOK_FUTURE (figure 3d) and the partial looking ahead algorithm is obtained. Forward checking consists of figure 3a and 3b with line 5 of figure 3a, the call to LOOK_FUTURE, deleted so that only CHECK_FORWARD is called. Forward checking does no checks of future units with future units.

```
1.  PROCEDURE LOOK_FUTURE(U,NEW_T);
2.  IF U+1 ≥ NUMBER_OF_UNITS THEN RETURN;
3.  FOR U1 = U+1 TO NUMBER_OF_UNITS BEGIN
4.      FOR L1 = each element of NEW_T(U1)
5.          FOR U2 = U+1 TO NUMBER_OF_UNITS except skipping U1 BEGIN
6.              FOR L2 = each element of NEW_T(U2)
7.                  IF RELATION(U1,L1,U2,L2) THEN
8.                      BREAK for L2 loop;    /* consistent label found */
9.              IF no consistent label was found for U2 THEN BEGIN
10.                  Delete L1 from list NEW_T(U1);
11.                  BREAK for U2 loop;        /* unit U2 has no label
                                                consistent with U1,L1 */
12.              END;
13.          END for U2 loop;
14.      END for L1 loop;
15.  IF NEW_T(U1) is empty THEN BEGIN
16.      NEW_T = EMPTY_ROW_FLAG;
17.      RETURN;
18.  END;
19.  END for U1 loop;
20.  RETURN;
21.  END LOOK_FUTURE;
```

Figure 3c

Figure 3c is the LOOK_FUTURE procedure, which deletes future labels which are not consistent with at least one label for every unit other than the label's own unit.

```
1.  PROCEDURE PARTIAL_LOOK_FUTURE(U,NEW_T);
2.  IF U+1 ≥ NUMBER_OF_UNITS THEN RETURN;
3.  FOR U1 = U+1 TO NUMBER_OF_UNITS - 1 BEGIN
4.      FOR L1 = each element of NEW_T(U1)
5.          FOR U2 = U1+1 TO NUMBER_OF_UNITS BEGIN
6.              FOR L2 = each element of NEW_T(U2)
7.                  IF RELATION(U1,L1,U2,U2) THEN
8.                      BREAK for L2 loop; /* consistent label found */
9.                  IF no consistent label was found for U2 THEN BEGIN
10.                     Delete L1 from list NEW_T(U1);
11.                     BREAK for U2 loop; /* unit U2 has no label
                                           consistent with U1,L1 */
12.                 END;
13.             END for U2 loop;
14.         END for L1 loop;
15.     IF NEW_T(U1) is empty THEN BEGIN
16.         NEW_T = EMPTY_ROW_FLAG;
17.         RETURN;
18.     END;
19. END for U1 loop;
20. RETURN;
21. END PARTIAL_LOOK_FUTURE;
```

Figure 3d

Figure 3d is the PARTIAL_LOOK_FUTURE procedure. It differs from LOOK_FUTURE (figure 3c) only at lines 1, 3, and 5. Each future unit-label pair is checked only with units in its own future.




Figure 4 illustrates a segment of a tree trace. Because backmarking remembers from the first visit to unit u which labels for u were compatible with the labels for units 1 through $v-1$ (segment A) do not have to be performed. Only those for units v through $u-1$ (segment B) have to be performed.

```
1.  RECURSIVE PROCEDURE BACKMARK(U,F,MARK,LOWUNIT);
2.  FOR F(U) = 1 TO NUMBER_OF_LABELS BEGIN
3.    IF MARK(U,F(U))  $\geq$  LOWUNIT(U) THEN BEGIN
4.      TESTFLAG = TRUE;
5.      I = LOWUNIT(U);
6.      WHILE (I < U) BEGIN      /* Find lowest failure */
7.        TESTFLAG = RELATION(I,F(I),U,F(U));
8.        IF NOT TESTFLAG THEN BREAK while loop;
9.        I = I + 1;
10.     END while loop;
11.     MARK(U,F(U)) = I;      /* Mark label with lowest failure */
12.     IF TESTFLAG THEN
13.       IF U < NUMBER_OF_UNITS THEN
14.         CALL BACKMARK(U+1,F,MARK,LOWUNIT);
15.       ELSE
16.         Output the labeling F;
17.       END;
18.   END for F loop;
19.   LOWUNIT(U) = U - 1;      /* Previous level will now change */
20.   FOR I = U+1 TO NUMBER_OF_UNITS;
21.     LOWUNIT(I) = MIN(LOWUNIT(I),U-1);
22.   RETURN;
23. END BACKMARK;
```

Figure 5 is Gaschnig's backmarking procedure as it was modified to find all solutions to constraint satisfaction problems (see Gaschnig, 1977).

III. Experimental Results

In this section we compare the six procedures, partial and full looking ahead, backtracking, backchecking, forward checking, and backmarking, on the N-queens problem for $4 \leq N \leq 10$, and on a random constraint problem. We assume that the unit order is fixed in its natural order from 1 to N and that all consistency tests of the current unit with past units or future units begin with the lowest ordered unit. The label sets will consist of all N columns; no consideration is given to the various symmetries peculiar to the N-queens problem.

The random constraint problems are generated using a pseudo-random number generator. A random number is generated for each possible consistency check $(u_1, l_1; u_2, l_2)$ for the relation R, so that each entry in the relation will be made with probability p. A probability of $p = 0.65$ is chosen so that problems will be generated that are somewhat similar to the N-queens problem.

The comparison among the tree search reducing procedures indicates that backtracking is least efficient in most cases, and that backmarking and forward checking are more efficient for the cases tested. Bit parallel forward checking, which takes advantage of machine parallelism, is the most efficient for all cases tried.

Our comparison of algorithm complexity will be in terms of nine criteria involving number of consistency tests, number of table lookups, and number of nodes in the tree search. There are a variety of ways of presenting these results including

1. Number of consistency tests performed to obtain all solutions (figures 6 and 8).

2. Number of table lookups used in finding all solutions (figures 10 and 11).
3. Number of nodes in the tree search to obtain all solutions (figure 12).
4. Number of nodes visited at each level in the tree search. (figure 13).
5. Number of nodes found to be consistent at each level in the tree search, or consistent labelings to depth (figures 17 and 18).
6. Number of consistency checks at each level in the tree search (figure 19).
7. Number of table lookups at each level in the tree search (figure 20).
8. Percentage of nodes at each depth that fail because an inconsistency was found at that depth (figures 21 and 22).
9. Average number of table lookups per consistency check (figures 23 and 24).

Figure 6 indicates that the number of consistency tests performed to obtain all solutions seems to increase exponentially with N for the N -queens problem. The number of solutions to the N -queens problem also appears to increase exponentially (see figure 7). The number of bit vector operations is also shown, for forward checking done in the bit vector data structure. Though backmarking appears to do slightly fewer consistency checks than forward checking on the N -queens problem, the use of machine parallelism gives bit parallel forward checking a clear advantage over all the other algorithms.

Random constraint problems with fixed probability .65 of consistency check success probability of 0.65 are tested in figures 8 and 9. The number of consistency tests appears to grow exponentially in figure 8, until a sufficiently large problem size is reached. At this point the number of solutions drops, as is indicated in figure 9, and the number of consistency tests appears to grow more slowly. Figure 9 explains the unevenness of the curves in figure 8. Too few random relations were tested for the means to settle closely to the expected values for this type of problem, and the average number of solutions varies erratically high and low.

In the random problems, forward checking does slightly fewer consistency checks than backmarking in the larger problem sizes, and once again matching parallelism gives bit parallel forward checking a clear advantage.

The number of table lookups for the N-queens and random relation problems are compared in figures 10 and 11. Only the lookups in the MARK array in backmarking and backchecking, and the T or NEW_T tables in the lookahead type algorithms are considered. These table lookups occur at line 3 in figure 5, backmarking, line 2 in figure 3a, line 4 in figure 3b, and lines 4 and 6 in each of figures 3c and 3d, the lookahead type algorithms. The entering of values into the tables are not considered, since they always follow at least one consistency check, and never happen more often than consistency checks.




Figure 6 compares the efficiency of the standard backtracking procedure with backchecking, looking ahead, forward checking, and backmarking, for the N-queens problem in the natural unit order. The number of bit vector operations for bit parallel forward checking is also shown.




Figure 7 illustrates the number of solutions for the N-queens problem.




Figure 8 represents the number of consistency tests in the average of 5 runs of the indicated programs. Relations are random with consistency check probability $p=0.65$ and number of units = number of labels = N . Each random relation is tested on all 6 methods, using the same 5 different relations generated for each N . The number of bit-vector operations in bit parallel forward checking is also shown for the same relations.




Figure 9 shows the average number of solutions for the probabilistic relations in figure 8. This is the average of 5 experiments for each problem size, and relations are random with consistency check success probability $p=0.65$. The dotted curve is the expected number of solutions.




Figure 10 shows the number of table lookups used to find all solutions to the N-queens problem for varying N, with the natural unit order.




Figure 11 shows the number of table lookups for the average of 5 random relations at each N, with number of units = number of labels = N, and probability of consistency check success $p = 0.65$. These random problems are the same as shown in figures 8 and 9.




Figure 12 illustrates the number of nodes in the tree search to find all solutions to the N-queens problem for varying N.




Figure 13 shows the number of nodes visited at each depth of the tree search to find all solutions of the 8-queens problem, for the various algorithms.

```
1 A
  2 C
  2 D
    3 B,F
  2 E,F
1 B
  2 D
    3 A
    3 F
      4 A
        5 C
          6 E
  2 E
    3 A,C
  2 F
```

figure 14 shows a segment of a tree trace made by the full looking ahead procedure in a 6-queens problem. One consistent labeling, 1B 2D 3F 4A 5C 6E, appears in this portion of the trace. 1A 2C fails to spawn any further nodes because the LOOK_FUTURE algorithm will, after deleting several potential labels, discover that one future unit has no possible labels.

```

1  A
  2  C
    3  E,F
  2  D
    3  B,F
  2  E
    3  B
  2  F
1  B
  2  D
    3  A
      4  C
    3  F
      4  A
        5  C
          6  E
  2  E
    3  A
      4  F
    3  C
  2  F

```

Figure 15 is a segment of a tree trace showing the nodes of the tree search in the partial looking ahead procedure in a 6-queens problem. One consistent labeling, 1B 2D 3F 4A 5C 6E, appears in this portion of the trace. 1A 2C 3A,B,C, and D do not appear because CHECK_FORWARD removes them from the table at nodes 1A and 2C. However 1A 2C 3E fails to have successors because the only labels left for future units 4 and 6 are incompatible and are removed by LOOK_FUTURE (see figure 1a).

```
1 A
  2 C
    3 E
      4 B
        3 F
          4 B
            2 D
              3 B
                4 E
                  3 F
                    4 C
                      2 E
                        3 B
                          4 F
                            5 C
                              2 F
                                3 B,D
                                  1 B
                                    2 D
                                      3 A
                                        4 C
                                          5 E
                                            3 F
                                              4 A
                                                5 C
                                                  6 E
                                                    5 E
                                                      4 C
```

Figure 16 shows a segment of a tree trace made by the forward checking procedure in a 6-queens problem. One consistent labeling, 1B 2D 3F 4A 5C 6E, is found. Notice that 1A 2C 3E 4B fails because the CHECK_FORWARD procedure discovers that there are no labels remaining for unit 6 at the 4B node (see figure 1b).




Figure 17 illustrates the number of consistent labelings as a function of tree depth, for the 8-queens problem with the natural unit order.

Backtracking is not shown, since it does no table lookups of the type considered. Partial and full looking ahead always do more table lookups than forward checking in these cases, and forward checking does better than backmarking in the larger problem sizes. Even full looking ahead does fewer table lookups than backmarking in the larger random problems. The number of table lookups into bit vectors is smaller than the number of table lookups in other algorithms, when the bit parallel data structure is used in forward checking.

Figure 12 demonstrates that full looking ahead visits the fewest nodes in the tree search, since it eliminates the most potential nodes during its examination of future unit-label pairs. Figure 13 indicates that the number of nodes visited in the tree search is largest for the middle levels in the tree search, with the full looking ahead procedure having the fewest nodes at each level.

Figures 14, 15, and 16 show segments of the trace of nodes visited by the full and partial looking ahead, and forward checking algorithms for the 6-queens problem. Backmarking and backchecking will have the same node trace as backtracking (see figure 2). More detailed trace of the action of backmarking can be found in Gaschnig, (1978 and 1979).

Figure 17 shows the number of consistent labelings at each depth of the tree search for the 8-queens problem, and figure 18 shows the average number of consistent labelings for random problems. This is the number of nodes at each level which have not yet been




Figure 18 illustrates the average number of nodes visited at each depth of the tree search for solutions to 5 random relations with number of units = number of labels = 10, and probability of consistency check success of $p = 0.65$. These are the 5 relations shown at the $N = 10$ case in figures 8, 9, and 11.




Figure 19 compares the number of consistency tests made at each level in the tree search for six different procedures, for the 8-queens problem, in the natural unit order.




Figure 20 shows the number of table lookups at each possible depth k in the tree search to find all solutions to the 8-queens problem, for the various algorithms.




Figure 21 shows the percentage of nodes at a given depth in the tree search which fail because some inconsistency is detected at that node. Results are shown for the various algorithms in the 8-queens problem. Backchecking and backmarking often discover that there is an inconsistency by using table lookups, rather than performing all the checks. Nodes in the lookahead type algorithms fail because a future unit fails to have any remaining labels after inconsistent future labels are removed.




Figure 22 shows the percentage of nodes at a given depth in the tree search which fail because some inconsistency is detected at that node. Results are shown for the average over 5 random relations, with consistency check probability $p=0.65$ and number of units = number of labels = 10.

found to be inconsistent. Backmarking and backchecking will have the same search tree as backtracking, and consequently has the same number of nodes and consistent labelings at each depth (see figures 13, 17, and 18). Their efficiencies are gained by reducing the ammount of work spent at each node, checking against past units. However, the lookahead algorithms perform extra work at each node to reduce the number of nodes, and as figures 19 and 20 show, the relation checks and table lookups for the lookahead type algorithms are concentrated more at the shallow depths of the tree search. As figures 6, 8, 10, and 11 show, full and partial looking ahead do too much work at each node for the problems shown, and forward checking and backmarking do better.

The percentage of nodes at each depth in the tree search that fail because some inconsistency is discovered are shown in figures 21 and 22. In the cases shown, in the backtracking, backchecking, and backmarking algorithms, over 95 percent of the nodes (instantiated labels) fail at the deepest level of the tree search, because they are inconsistent with some past unit-label pair.

The lookahead type algorithms reduce the number of nodes in the tree search in two ways. First by removing entries from the tables of potential unit-label pairs, and second by noticing that some future units may have no possible labels associated with them. In Figures 21 and 22 all nodes that fail in the lookahead type algorithms do so for this second reason, since they would not have occurred as nodes if their labels were deleted from the future unit-label tables. If lines 7 and 8 are removed from the CHECK_




Figure 23 shows that in the N-queens problem, all the algorithms that keep tables, except backmarking, reference those tables no more than 35 percent more often than they reference the relation. However backmarking appears to have a continually growing ratio of table lookups to relation tests, as problem size grows.




Figure 24 demonstrates that in the average over 5 random relations for each problem size that was tested, that the ratio of table lookups to consistency checks seems to approach one, except in the backmarking algorithm, in which the ratio seems to steadily increase with problem size. The random relations are the same as those shown in figures 8, 9, and 11.

FORWARD procedure in figure 3b, then no nodes would fail in forward checking (this is McGregor's restricted arc consistency algorithm in McGregor, 1979). This weaker form of forward checking algorithm will find all the consistent labelings at each level of the tree search that backtracking does, but at a higher cost than the original forward checking algorithm. The replacement of these lines in forward checking will realize a 15 percent saving of consistency checks and table lookups in the 10-queens problem, and over 40 percent savings in the 10 units by 10 labels random problem.

Figures 23 and 24 address the question of what measure best determines the algorithmic time complexity. A careful check of all the algorithms will show that no step is executed more often than the maximum of the number of consistency checks or the number of table lookups. As the problem size increases in the lookahead type of algorithms, the ratio of table lookups to consistency checks seems to decrease from a maximum of about 2 to no more than 1.5 table lookups per consistency check in both the N-queens and random problems. This ratio is guaranteed to be greater than or equal to one, by the algorithms structure, thus either may be used as a measure of algorithmic time complexity for the lookahead type algorithms.

Because in both the N-queens and random problems backmarking seems to have a steadily increasing ratio of table lookups to consistency checks as problem size grows, only table lookups (which equals the number of nodes in the tree search in this case) can be used as a true measure of algorithmic time complexity for backmarking.

Only in the case that a computation of a relation check is significantly more expensive than the cost of a node's loop control and a table lookup will relation checks be a useful practical measure for the time complexity of backmarking. The reason that it is a practical measure in this case is that the node and table lookups cost will dominate the cost of execution only in very large problem sizes, so large that the problems can not be solved in a reasonable time, and relation tests will dominate the cost in the smaller problems which can be solved in a practical amount of time.

IV. Statistical Model for Constraint Satisfaction Searches

Our statistical model for random constraint satisfaction is simple. The probability that a given consistency check succeeds is independent of the pair of units or labels involved and is independent of whatever labels may already have been assigned to past units. Hence, $P((u_{k+1}, l_{k+1}, u, l) \in R | l_1, \dots, l_k \text{ are consistent labels of } u_1, \dots, u_k) = P((u_{k+1}, l_{k+1}, u, l) \in R)$ for every u, l .

The N-queens problem is a more difficult problem, with fewer solutions but requiring more consistency tests than the corresponding random constraint problem with the same probability of consistency check success. A comparison of the graphs for the two problems in section III will show that while the numerical values of the quantities vary considerably, the basic character of the algorithms operation is similar for both problems.

In our analysis, we will assume that a given pair of units with a given pair of labels is consistent with probability p , p being independent of which units, which labels, or any past processing. If each unit has the same number, M , of possible labels, then any k -tuple of labels for any k units has probability $p^{k(k-1)/2}$ of being consistent since each labeling must satisfy $k(k-1)/2$ consistency checks. Since there are M^k possible labelings of k units, the expected number of consistent labelings is

$$M^k p^{k(k-1)/2}.$$

The expected number of nodes processed at level k in a standard backtracking search will be M , the number of possible labels, times the number of consistent labelings at the previous level,

$M^{k-1} p^{(k-1)(k-2)/2}$. Thus there are

$$M^k p^{(k-1)(k-2)/2}$$

tree search nodes at level k .

We can also count the expected number of consistency checks performed by backtracking. We expect $M^k p^{(k-1)(k-2)/2}$ level k nodes and at each node a label must be tested for consistency with the labels given the previous $k-1$ units. The first consistency check fails with probability $1-p$. If it fails, we have spent 1 test. If it succeeds we have spent 1 test and are committed to make another one which might also succeed with probability p . All $(k-1)$ tests will succeed with probability $p^{(k-1)}$. Hence the expected number of consistency checks performed at each node is

$$\sum_{i=1}^{k-1} i p^{i-1} (1-p) + (k-1) p^{k-1}.$$

This may be simplified by recognizing the telescopic nature of the sum which is equal to

$$\sum_{i=0}^{k-2} p^i.$$

But this is a geometric sum and is equal to

$$\frac{1 - p^{k-1}}{1 - p}.$$

Therefore the expected number of consistency checks at level k will be $M^k p^{(k-1)(k-2)/2}$, the number of nodes at level k times $\frac{1 - p^{k-1}}{1 - p}$, the expected number of consistency checks at a node,

making

$$M^k p^{(k-1)(k-2)/2} \frac{1 - p^{k-1}}{1 - p}$$

consistency checks at level k . Of course the expected total number of consistency checks will be the summation of the expected number of consistency checks for each level k for k ranging from 1 to N , the number of units.

The computation of the number of labelings for the forward checking algorithm is somewhat more complicated because the algorithm stops checking when a future unit has no labels that are consistent with the past and present unit-label pairs. A consistent labeling to depth k occurs when the tree search successfully reaches a given label for unit k and forward checking of that unit-label pair produces no future unit that has no remaining labels. Thus the consistent labelings to depth k for forward checking meet exactly the following conditions:

- 1) $u_1\ell_1, u_2\ell_2, \dots, u_k\ell_k$ are consistent unit-label pairs.
- 2) There is no future unit u in levels $k+1, \dots, N$ for which there is no label ℓ so that u, ℓ is consistent with $u_1\ell_1, \dots, u_k\ell_k$.

The k unit-label pairs are consistent with probability $p^{k(k-1)/2}$, and there are M^k possible labelings to depth k (condition 1). A future unit-label pair is consistent with the k past and present unit-label pairs with probability p^k and there are M possible labels for a future unit, so the probability that a future unit has no label that is consistent with the k past and present units is $(1 - p^k)^M$. Since there are N units, there are $N-k$ future units,

and the probability that all of these has at least one label that is consistent is $[1 - (1-p^k)^M]^{N-k}$ (condition 2). Thus the expected number of consistent labelings to depth k for forward checking is

$$M^k p^{k(k-1)/2} [1 - (1-p^k)^M]^{N-k}.$$

The expression for the expected number of nodes in the forward checking tree search at level k is very similar to that for the number of consistent labelings to depth k, since each node will perform forward checks to determine if its label will become a consistent labeling. The labels for a node must meet condition 1 above, but the future units are required to have succeeded with at least one label only for checks with the unit label pairs $u_1 \ell_1, \dots, u_{k-1} \ell_{k-1}$, since each node was in the table for a consistent labeling to depth k-1, and if any future unit as seen from level k-1 failed to have a label then it would not have spawned nodes at the next level. Thus in each node the future units will have at least one label and the second condition occurs with probability $[1 - (1-p^{k-1})^M]^{N-k}$. Thus the expected number of nodes at depth k in forward checking is

$$M^k p^{k(k-1)/2} [1 - (1-p^{k-1})^M]^{N-k}.$$

A slight overapproximation for the expected number of consistency checks at depth k in the tree search can be found by multiplying the expected number of nodes at the depth times the expected number of labels remaining for each future unit times the number of future units, N-k. Since each future unit will have at least one label, this expected number of labels will be $\frac{M p^{k-1}}{1 - (1-p^{k-1})^M}$. Thus the expected

number of consistency checks in forward checking will be

$$M^{k+1} p^{(k+2)(k-1)/2} [1 - (1-p^{k-1})^M]^{N-k-1} (N-k).$$

The exact expected value can be obtained by replacing the number of future units term, $(N-k)$, with the expected number of future units tested, since forward checking will stop testing as soon as a future unit is discovered to have no possible labels. Each of these tests of a future unit will succeed with probability

$1 - (1-p^k)^M$, and reasoning similar to that for the number of consistency checks at each node in backtracking will give

$$\frac{1 - [1 - (1-p^k)^M]^{N-k}}{(1-p^k)^M} \text{ for the expected number of future units}$$

tested. Thus the expected number of consistency checks at level k in forward checking will be

$$M^{k+1} p^{(k+2)(k-1)/2} [1 - (1-p^{k-1})^M]^{N-k-1} \cdot \frac{1 - [1 - (1-p^k)^M]^{N-k}}{(1-p^k)^M}.$$

The number of table lookups in forward checking is the sum of the number of consistency checks and the number of nodes. Thus the expected number of table lookups at depth k in the tree search will be

$$M^k p^{k(k-1)/2} [1 - (1-p^{k-1})^M]^{N-k} \left[1 + \frac{M p^{k-1}}{1 - (1-p^{k-1})^M} \cdot \frac{1 - [1 - (1-p^k)^M]^{N-k}}{(1-p^k)^M} \right].$$

The expected number of bit vector operations with the bit parallel data structure in forward checking can easily be found, by removing the term for the number of labels remaining for each future unit from the expression for the number of consistency tests, since only one operation will be performed for each unit, giving

$$M^k p^{k(k-1)/2} [1 - (1-p)^{k-1} M]^{N-k} \frac{1 - [1 - (1-p)^k M]^{N-k}}{(1-p)^k M}$$

for the expected number of bit vector operations at level k in bit parallel forward checking.

The number of table lookups in bit vectors is still the sum of the number of bit vector operations and the number of nodes, for

$$M^k p^{k(k-1)/2} [1 - (1-p)^{k-1} M]^{N-k} \left[1 + \frac{1 - [1 - (1-p)^k M]^{N-k}}{(1-p)^k M} \right]$$

table lookups into bit vectors in bit parallel forward checking.

To illustrate the general form of the expressions we computed for the expected number of consistency checks and expected number of solutions, we present a few graphs. Figure 25 illustrates the graph of the expected number of consistency checks as a function of tree depth for a random constraint satisfaction problem having $N = 17$ units and labels and a probability $p = 0.70$ of a constraint being satisfied. Notice that the greater number of tests forward checking does early in the tree search pays off later in the tree search both in terms of number of consistency tests and in number of successful instantiations at each tree depth (figure 26).

Figure 27 illustrates the expected number of solutions as a function of N and p parameters of a random constraint satisfaction problem. Increasing N for a fixed p eventually causes fewer solutions to exist because the number of constraints is increasing quadratically.

We, of course, expect the number of consistency tests to increase as N increases and p remains fixed since the search space is becoming

large very rapidly. This is shown for the forward checking procedure in figure 28. Also expected is for the average number of consistency checks per labeling to increase as N increases and p remains fixed (figure 29). As N increases, the problem of finding the first solution as well as all solutions is becoming more and more difficult. Therefore, it is not expected for the number of consistency tests per solution to decrease as the number of solutions increases. The reason for this is that as the number of solutions increases more of the tests required to verify a solution become shared because the solutions have common segments. This is illustrated in figure 30.

An experimental check of the theoretical equations for the number of solutions and number of solutions at a given depth in the tree search with random relations is given in figures 31 and 32. Although the average number of solutions is close to the theoretical result, the individual relations vary widely. The I bars mark a distance of one standard deviation of the mean above and below the average of the trials.

Figure 33 demonstrates the accuracy of the theoretical expression for the expected number of consistency tests with random relations in the forward checking algorithm. The total expected number of consistency tests shown is calculated from the sum of the expected number of tests for each level in the various problem sizes. For this expression to be correct, the expression for the expected number of nodes at each level in the forward checking algorithm must also be correct.




Figure 25 illustrates the number of consistency tests as a function of tree depth for an $N = 17$, $p = 0.70$ random constraint satisfaction problem.




Figure 26 illustrates the number of consistent labelings as a function of tree depth for an $N = 17$, $p = 0.70$ random constraint satisfaction problem.




Figure 27 illustrates the expected number of solutions as the N and p parameters change for a random constraint satisfaction problem.




Figure 28 shows how the number of consistency tests increases as N increases, and p is held constant.




Figure 29 shows how the number of consistency tests per solution increases as N increases and p is held constant, in forward checking.




Figure 30 illustrates that as the expected number of solutions increases, the expected number of consistency checks per solution decreases, for varying probability p and fixed problem size N = number of units = number of labels, in the forward checking algorithm.




Figure 31 shows the expected number of consistent labelings for random relations with number of units = number of labels = N , and consistency check success probability $p = 0.65$. Dots represent the average number of consistent labelings for 25 random relations tested at each problem size, and the I bars indicate one standard deviation of the mean above and below the experimental mean.




Figure 32 indicates the number of consistent labelings to depth k in the tree search for the average of 25 random constraint satisfaction problems with probability of consistency check success of 0.65 and number of units = number of labels = 10. The dotted curve is the theoretical expected number of labelings for such a problem.




Figure 33 shows the expected number of consistency checks for random relations with number of units = number of labels = N , and consistency check success probability $p = 0.65$. Dots represent the average number of consistency checks for 25 random relations tested at each problem size, and the I bars indicate one standard deviation of the mean above and below the experimental mean.

V. The Fail First Principle

One of the strategies which helps tree searching for constraint satisfaction problems is the fail first or prune early strategy of the looking ahead and forward checking procedures. There are other ways that we can apply the general principle of trying to fail first (and of course remember that fact so that there are no unnecessarily repeated mistakes). In this section we discuss two other applications of this strategy. The first is by optimizing the order in which we do consistency tests. The fail first principle states that we should first try those tests in the given set of tests that are most likely to fail since if they do fail we do not have to do the remainder of the tests in the set.

The second application is in dynamically choosing the optimal order in which to process units in each branch of the tree search. Optimal unit order choosing, even on a local basis, will not only lower the number of expected consistency tests per problem as compared with a random ordering, but it also lowers the variance of this average. For the unit order choice, the fail first principle states that the next unit to choose should be that one with the fewest possible labels left.

V.1 Optimizing the Consistency Check Order in Tree Searching

Suppose we are solving a constraint satisfaction problem and suppose units $1, \dots, K$ have already been assigned labels $\ell_1, \dots, \ell_K$

and we are trying to find a label ℓ_{K+1} for unit $K+1$. The label ℓ_{K+1} must come from some set S_{K+1} of labels and it must be consistent with each of the previous labels $\ell_1, \dots, \ell_K$, that is, we must have $(k, \ell_k, K+1, \ell_{K+1}) \in R$ for $k = 1, \dots, K$. To determine the label ℓ_{K+1} , we sequentially go through all the labels in S_{K+1} and perform the K consistency checks: $(k, \ell_k, K+1, \ell_{K+1}) \in R$. If one check fails, then we try the next label in S_{K+1} . If all checks succeed, then we can continue the depth first search with the next unit.

The optimizing problem for consistency checking is to determine an order in which to perform the tests which minimizes the expected number of tests performed. To set up the optimizing problem, we must have some knowledge about the degree to which a previous unit's label constrains unit $(K+1)$'s label. For this purpose we let $P(k)$ be the probability that the label ℓ_k for unit k is consistent with some label for unit $K+1$. We assume that the consistency checks are independent events so that the probability of the tests succeeding on units 1 through K is $\prod_{k=1}^K P(k)$.

For each order of testing, these probabilities determine the expected number of tests in the following way. Let $k_1, \dots, k_K$ be a permutation of $1, \dots, K$ designating the order in which the consistency checks will be performed. The test $(k_1, \ell_{k_1}, K+1, \ell_{K+1}) \in R$ will succeed with probability $P(k_1)$ and fail with probability $1-P(k_1)$. If it fails, we incur a cost of one consistency check and we try the next label. If it succeeds, we will have incurred

a cost of one consistency check and we are committed to try the next test $(k_2, \ell_{k_2}, K+1, \ell_{K+1}) \in R$. This test succeeds with probability $P(k_2)$ and fails with probability $1-P(k_2)$. At this point, we have incurred a cost of two tests and may be committed to make more tests if this one succeeded.

Figure 34 shows the tree of $K+1$ possible outcomes. Since the tests are assumed independent, the probability for each outcome can be computed by multiplying probabilities. For example, the probability of failing on consistency check with unit k_3 is $P(k_1)P(k_2)(1-P(k_3))$. Also associated with each outcome is the number of tests performed to get there. For example, failure on the test with unit k_3 incurs a cost of 3 tests.

The expected number of tests C performed is computed by

$$C = \sum_{i=1}^K i [1-P(k_i)] \prod_{j=1}^{i-1} P(k_j) + \sum_{i=1}^K \prod_{j=1}^i P(k_j)$$

Upon rearranging and simplifying this expression we obtain

$$C = 1 + \sum_{i=1}^{K-1} \prod_{j=1}^i P(k_j)$$

Now by the proposition at the end of section V, this is minimized by having $k_1, \dots, k_K$ be any permutation of $1, \dots, K$ satisfying $P(k_1) \leq P(k_2) \leq \dots \leq P(k_K)$. Hence, to minimize expected numbers of tests, we must choose the order so that the tests with units most likely to fail are done first.

To illustrate the advantage of using optimum consistency test order, we consider the 10-queens problem when the units are naturally ordered from 1 to N and the current unit is K , then the fail

first principle states that tests with past units must be done in the order of decreasing constraints. Since the row previous to row k has the strongest constraint on row k , the test order should be first unit $K-1$, then $K-2$, up to unit 1, in the N -queens problem. Backtracking requires 1,297,488 tests when done in the wrong order (unit 1, 2, ..., $K-1$) and 1,091,856 tests when done in the right order. It is interesting to note that Gaschnig's backjumping procedure (Gaschnig, 1978) when done with the consistency tests in the wrong order (1,131,942 tests) performs worse than standard backtracking with consistency tests in the right order. Furthermore, for the N -queens problem backjumping with consistency tests in the right order for the N -queens problem is equivalent to standard backtracking with consistency tests in the right order because backjumping backtracks to the highest level at which a failure is detected, and there is always at least one table at a given level which fails when checked with the immediately preceding level.




Figure 34 illustrates the $K+1$ outcomes of K tests.

V.2 Optimizing Tree Search Order

Every tree search must assume some order for the units to be searched in. The order may be uniform through the tree or may vary from branch to branch. It is clear from experimental results that changing the search order can influence the average efficiency of the search. In this section we adopt the efficiency criterion of branch depth and we show how by always choosing the next unit having smallest number of label choices we can minimize the expected branch depth.

Suppose units $1, \dots, N$ are units which are yet to be assigned labels. Let $n(m)$ be the number of possible or available labels for unit m . We assume that each of the $n(m)$ labels possible for unit m has the same probability q of succeeding and that success or failure of one of the labels is an independent event from success or failure for any of the other labels. Thus, the probability that a unit m will not have any label that succeeds is $(1 - q)^{n(m)}$. The probability that some label for unit m succeeds is, therefore, $P(m) = 1 - (1 - q)^{n(m)}$. Unfortunately, this analysis holds only for the first level of the tree.

Let $k_1, \dots, k_N$ be the order in which the units are searched on the tree. Let $P_n(k_n | k_1, \dots, k_{n-1})$ be the conditional probability that some label for unit k_n will succeed when unit k_n is the n^{th} one in the tree search order given that units $k_1, \dots, k_{n-1}$ are the first $n-1$ units searched in the branch. We assume that the probability of a label for unit k_n succeeding depends only on the number

of units preceding it in the tree search and not upon which particular units they are. That is,

$$P_n(k_n | k_1, \dots, k_{n-1}) = P_n(k_n | \ell_1, \dots, \ell_{n-1})$$

for all labels $\ell_1, \dots, \ell_{n-1}$. This conditional independence assumption justifies the use of the notation $P_n(k_n)$ to designate the probability that some label succeeds for unit k_n when it is the n^{th} unit in the tree search, and we will call the probability that an arbitrary label for unit u will succeed when checked against another arbitrary unit-label pair the success probability for unit u .

Units which are searched later in the tree typically have lower probability for a label succeeding since the label must be consistent with the labels given all the earlier units. We want some way to compare the probability of success for the same unit in different tree searches. Since the success probability depends only on the unit and its level in the tree and since units later in the tree have lower success probabilities, we assume that the success probability for a unit u when it is at level i in one tree search is related to the success probability of unit u when it is at the first level of another tree search by a constant factor α^{i-1} where $0 < \alpha \leq 1$:

$$P(u) \alpha^{i-1} = P_i(u).$$

The best search order is the one which minimizes the expected length or depth of any branch. When the units are searched in

the order $k_1, \dots, k_m$, the expected branch depth is given by

$$1 + \sum_{n=1}^{N-1} \prod_{j=1}^n P_j(k_j).$$

By the proposition at the end of section V, this is minimized when the unit chosen at each level is that unit whose success probability is smallest. Thus at level j we choose unit k_j , where

$$P_j(k_j) \leq P_j(u) \quad \text{for } u \neq k_1, \dots, k_{j-1}.$$

Now, $P_j(k_j) = \alpha^{j-1} [1 - (1 - q^{n(k_j)})]$. Since $0 \leq q \leq 1$, this expression is minimized by choosing k_j to be that unit having the smallest number of possible labels.

To illustrate the advantage of using a locally optimal unit order for each branch in the tree search, we consider the improvement achieved on the N -queens problem and random relation problems. The number of consistency tests required is given in tables 1 and 2. Some improvement is shown in the larger N -queens problems, and considerable improvement appears in the larger random relation problems. Figure 35 demonstrates that the improvement increases with problem size in the random relation problems with $p=0.65$ and number of units = number of labels = N .

The reason why optimal unit order usually improves forward checking more than backmarking is that forward checking has more information about future units than backmarking. Therefore, forward checking's choice of the next unit most likely to fail is more likely to produce a unit which fails than backmarking's choice.

Table 1 Number of Consistency Tests in N-queens
Problem for Normal and Optimal Unit Order

N	Backtracking	Backchecking	Full Looking Ahead		Part Looking Ahead		Backmarking		Forward Checking	
	Normal	Normal	Normal	Optimal	Normal	Optimal	Normal	Optimal	Normal	Optimal
4	84	80	99	99	97	97	76	76	76	76
5	405	356	598	578	485	431	276	276	282	282
6	2016	1496	2095	2082	1703	1708	944	921	964	946
7	9297	6042	8942	8941	6511	6318	3236	3168	3338	3229
8	46752	27450	35323	35211	25882	25062	12308	12095	13024	12108
9	243009	131538	153455	151275	112327	106247	50866	50027	55326	49856
10	1297558	643658	661017	636377	496455	449666	220052	211635	242174	205970

Table 2 Number of Consistency Tests in Average of 5 Random Constraint
Satisfaction Problems with Consistency Check Success Probability
0.65, for Normal and Optimal Unit Order

N = number of units = number of labels

N	Backtracking	Backchecking	Full Looking Ahead		Part Looking Ahead		Backmarking		Forward Checking	
	Normal	Normal	Normal	Optimal	Normal	Optimal	Normal	Optimal	Normal	Optimal
4	243	195	230	223	184	174	132	133	133	121
5	1043	731	760	712	599	554	425	382	414	358
6	4637	2514	2543	2288	1966	1700	1298	1159	1273	979
7	12040	6722	5722	5156	4697	3875	3175	2486	3057	2069
8	25893	13490	10779	9306	9111	7233	6089	4300	5978	3425
9	118086	55318	30799	25904	26788	19174	21170	10246	18616	8673
10	163983	73260	44655	36675	41232	28872	28314	14892	25298	12022




Figure 35 shows the improvement in efficiency for the average of 5 random relation problems with probability $p=0.65$ and number of units = number of labels = N , when locally optimized unit order is used. The relations are the same as those shown in figures 9 and 10.

Proposition: Let $0 < \alpha \leq 1$ be given. For each unit u let $P(u)$ be its initial success probability. Let $k_1, \dots, k_N$ be any permutation of $1, \dots, N$ satisfying $P(k_1) < P(k_2) < \dots < P(k_N)$. Define $P_n(u) = \alpha^{n-1} P(u)$. Then,

$$\sum_{n=1}^{N-1} \prod_{j=1}^n P_j(k_j) \leq \sum_{n=1}^{N-1} \prod_{j=1}^n P_j(u_j) \text{ for any permutation } u_1, \dots, u_N \text{ of } 1, \dots, N.$$

Proof: Let $u_1, \dots, u_N$ be any permutation of $1, \dots, N$ minimizing

$$\sum_{n=1}^{N-1} \prod_{j=1}^n P_j(u_j).$$

If $u_1, \dots, u_N$ we are done. If $u_1, \dots, u_N$ does not

equal $k_1, \dots, k_N$, let m be the smallest index such that $u_m \neq k_m$.

Also let m' be the index such that $u_m \neq k_{m'}$. Also let m' be the

index such that $u_{m'} = k_m$. Define the permutation $i_1, \dots, i_N$ by

$$i_n = u_n, \quad n \neq m \text{ or } m'$$

$$i_m = u_{m'}$$

$$i_{m'} = u_m$$

We will prove a contradiction by showing that

$$\sum_{n=1}^{N-1} \prod_{j=1}^n P_j(i_j) < \sum_{n=1}^{N-1} \prod_{j=1}^n P_j(u_j)$$

by looking at the products. There are three cases: $n < m$, $m \leq n < m'$,

and $m' \leq n$.

Case 1 $n < m$: Here since $i_j = u_j$, $j=1, \dots, n$, we obtain

$$\prod_{j=1}^n P_j(i_j) = \prod_{j=1}^n P_j(u_j).$$

Case 2 $m \leq n < m'$:

$$\prod_{j=1}^n P_j(i_j) = \frac{P_m(u_m)}{P_m(u_m)} \prod_{j=1}^n P_j(i_j) = \frac{P_m(u_{m'})}{P_m(u_m)} \prod_{j=1}^n P_j(u_j)$$

Now, $P_m(u_{m'}) = P_m(k_m) < P_m(u)$ for any $u \neq k_1, \dots, k_{m-1}$. Since $u_m \neq k_1, \dots, k_{m-1}$, $P_m(u_{m'}) < P_m(u_m)$ and

$$\prod_{j=1}^n P_j(i_j) = \frac{P_m(u_{m'})}{P_m(u_m)} \prod_{j=1}^n P_j(u_j) < \prod_{j=1}^n P_j(u_j)$$

Case 3 $m' \leq n$:

$$\begin{aligned} \prod_{j=1}^n P_j(i_j) &= \frac{P_m(u_m)P_{m'}(u_{m'})}{P_m(u_m)P_{m'}(u_{m'})} \prod_{j=1}^n P_j(i_j) \\ &= \frac{P_m(i_m)P_{m'}(i_{m'})}{P_m(u_m)P_{m'}(u_{m'})} \prod_{j=1}^n P_j(u_j) \\ &= \frac{P_m(u_{m'})P_{m'}(u_m)}{P_m(u_m)P_{m'}(u_{m'})} \prod_{j=1}^n P_j(u_j) \\ &= \frac{\alpha^{m-1}P(u_{m'})\alpha^{m'-1}P(u_m)}{\alpha^{m-1}P(u_m)\alpha^{m'-1}P(u_{m'})} \prod_{j=1}^n P_j(u_j) \\ &= \prod_{j=1}^n P_j(u_j) \end{aligned}$$

Hence, $\sum_{n=1}^{N-1} \prod_{j=1}^n P_j(u_j) < \sum_{n=1}^{N-1} \prod_{j=1}^n P_j(u_j)$, contradicting the minimality

of $u_1, \dots, u_N$. Therefore, $u_1, \dots, u_N = k_1, \dots, k_N$.

VI. Conclusion

Using complexity criteria of number of consistency checks and number of table lookups we have shown analytically and experimentally the efficacy of the remembering and fail first principles in constraint satisfaction tree search problems. A new search procedure called forward checking has been described and it combined with optimal unit order choice leads to a more efficient tree search than looking ahead or backmarking. A data structure that takes advantage of a computer's natural ability to process bit vectors in parallel can make forward checking even more efficient. This suggests that the entire set of look ahead operators described by Haralick et. al. (1978) Haralick and Shapiro (1979a, 1979b), the discrete relaxation described by Waltz (1972) and Rosenfeld et. al. (1976) would be more efficiently implemented by omitting the consistency tests required by future units against future units. Further analytic and experimental work needs to be done to determine if this in fact is generally true. Applicability of the forward checking idea to inference and theorem proving algorithms needs to be tested and this will be the topic of a future paper.

REFERENCES

1. Barrow, H. G. and J. M. Tenenbaum, "MYSIS: A System for Reasoning about Scenes." Stanford Research Institute, Menlo Park, CA, SRI AI Tech. Rep. 121, 1976.
2. Gaschnig, J., "A General Backtrack Algorithm that Eliminates Most Redundant Tests." Proceedings of the International Conference on Artificial Intelligence, Cambridge, MA, 1977, p. 457.
3. Gaschnig, J., "Experimental Case Studies of Backtrack vs. Waltz-type vs. New Algorithms for Satisficing-Assignment Problems." Proceedings of the 2nd National Conference of the Canadian Society for Computational Studies of Intelligence, Toronto, Ontario, Canada, July 19-21, 1978.
4. Gaschnig, J., Performance Measurement and Analysis of Certain Search Algorithms, thesis, Department of Computer Science, Carnegie-Mellon University, May 1979.
5. Haralick, R. M., L. S. Davis, A. Rosenfeld, and D. L. Milgram, "Reduction Operations for Constraint Satisfaction." Information Sciences, Vol. 14, 1978, p. 199-219.
6. Haralick, R. M., "Scene Analysis: Homomorphisms, and Arrangements." Machine Vision, Hanson and Reisman, Eds., New York: Academic, 1978.
7. Haralick, R. and L. Shapiro, "The Consistent Labeling Problem: Part I." IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 1, No. 2, April 1979.
8. Haralick, R. and L. Shapiro, "The Consistent Labeling Problem: Part II." IEEE Transactions on Pattern Analysis and Machine Intelligence.
9. Harary, F., Graph Theory, Reading, MA: Addison-Wesley, 1969.
10. Kowalski, R., "A Proof Procedure Using Connection Graphs." J. Agg. Comput. Mach., Vol. 22, p. 572-595, Oct. 1975.
11. Mackworth, A. K. "Consistency in Networks of Relations." Artificial Intelligence, Vol. 8, No. 1, 1977, p. 99-118.
12. McGregor, J. J., "Relational Consistency Algorithms and their Applications in Finding Subgraph and Graph Isomorphisms." Information Sciences, Vol. 18, 1979.
13. Montanari, V., "Networks of Constraints: Fundamental Properties and Applications to Picture Processing." Information Sciences, Vol. 7, 1976, p. 95-132.

14. Rosenfeld, A., R. Hummel, and S. Zucker, "Scene Labeling by Relaxation Operations." IEEE Transactions on Systems, Man and Cybernetics, SMC-6, 1976, p. 420-433.
15. Waltz, D. L., "Generating Semantic Descriptions from Drawings of Scenes with Shadows." Report MAC AI-TR-271 MIT, Cambridge, Massachusetts, 1972.
16. Ullman, J. R., "Associating Parts of Patterns." Information and Control, Vol. 9, No. 6, December 1966, p. 583-601.

IED
80