Expensive Constraints and HyperArc-Consistency

Rajiv Chopra, Rohini Srihari and Anthony Ralston
Dept. of Computer Science, SUNY at Buffalo
224 Bell Hall, Amherst NY 14214
{rchopra,rohini,ralston}@cs.buffalo.edu

Abstract

We present an arc-consistency algorithm NAC4 to handle some types of non-binary constraints. It differs from the generalized algorithm GAC4 in that the constraints are not modeled as explicit relations but are retained as implicit n-ary relations. We believe that when the checking of constraints is expensive, NAC4 will perform better than GAC4. In NAC4 constraints are ordered according to their cost and cheaper constraints are applied first. It is reasonable to expect that the cost of checking the expensive constraints is lowered because of the reduction in the domains of some of the nodes as a result of checking the cheaper constraints first. This algorithm is suited for non-binary constraints between a node and a set of other nodes.

1 Introduction

Consistency algorithms, also called discrete relaxation algorithms, were first introduced by Waltz for interpreting polyhedral scenes with shadows [Waltz, 1975] posed as a labeling problem. Traditional techniques for constraint satisfaction include chronological backtracking and backtracking with forward checking [Haralick and Elliot, 1980]. The latter prunes the search tree of nodes based upon incompatibility with the partial labeling constructed so far. Consistency algorithms go one step further in removing incompatible future labelings. These incompatible labelings would have otherwise been discovered repeatedly by most backtracking procedures. Discrete relaxation runs in polynomial time to produce a locally consistent solution. When the Constraint Satisfaction Problem (CSP) is tightly constrained, application of consistency algorithms can greatly reduce or totally eliminate subsequent backtracking. These techniques have found wide application in artificial intelligence, pattern recognition and image interpretation.

The labeling problem can also be solved using Continuous relaxation techniques [Faugeras and Berthod, 1981], that are capable of handling fuzzy constraints, are also used for solving the labeling problem. However, these are difficult to control because the user has to juggle various parameters of the system to obtain the correct solution. In contrast, discrete relaxation techniques handle only boolean constraints and are easier to manipulate.

Many algorithms for binary acr-consistency have been proposed: **AC3** [Mackworth and Freuder, 1985], **AC4** [Mohr and Henderson, 1986] and **AC5** [Hentenryck *et al.*, 1992]. A generalization of AC4 to handle non-binary constraints, **GAC4**, was presented in [Mohr and Masini, 1988].

Consistency algorithms for finite domains assume that constraints are enumerated as explicit relations with little or no cost, and that verifying the existence of a tuple in a relation is a unit cost operation. However, there are applications of CSP, in domains such as computer vision [Chopra and Srihari, 1995], where enumerating the complete extension of a constraint relation is neither inexpensive nor necessary. When dealing with expensive constraints, reducing the domains of the relation's nodes prior to computing the relation's extension seems an obvious improvement over computing the extension first.

The rest of the paper is organised as follows. Section 2 provide basic definitions relevant to this work. Section 3 describes the optimal arc-consistency algorithms. Section 4 presents our new algorithm, NAC4, for non-binary constraints. Section 5 briefly analyses NAC4. Section 6 considers expensive constraints. Section 7 discusses an application of the algorithm.

2 Definitions

A *static* constraint satisfaction problem [Mackworth, 1977] is defined as a 4-tuple (V, D, C, R) where V is a set of **nodes**, $\{v_1, v_2, \ldots v_m\}$, each with an associated **domain** D_i of possible **labels**. (In Fig.1, adapted from [Bessiere, 1991], D_2 is {b,c,d})

The search space D, consists of the Cartesian

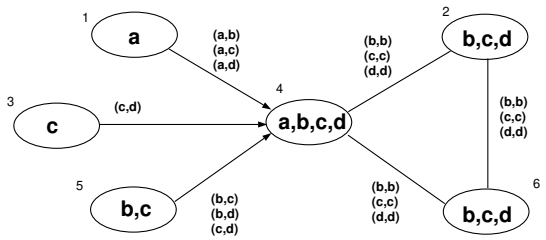

Figure 1: An example of a binary CSP with constraints modeled as explicit relations. The ovals are nodes, values in them are labels and the lists of pairs are constraints.

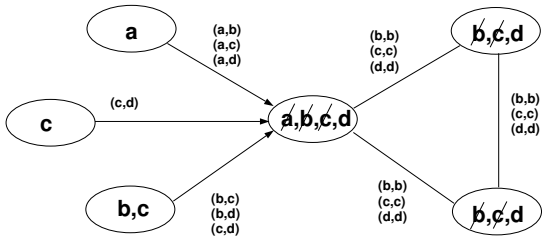

Figure 2: The CSP of fig.1 after application of an arc-consistency algorithm

product of the nodes' domains, that is $D_1 \times D_2 \times \ldots \times D_n$.

A set of **constraints** C is specified over some subsets of these nodes, where each constraint C_p involves a subset $\{vi_1, \ldots, vi_q\}$ of V. C_p is labeled by a relation R_p of R, a subset of the Cartesian product $D_{i_1} \times \ldots \times D_{i_q}$ that specifies those labels of the nodes that are compatible with each other.

A **binary** constraint involves only two nodes. This notion is extensible to an **n-ary** constraint over n nodes. The task in a CSP is to assign labels to nodes such that all the constraints on the nodes are satisfied simultaneously. Problem reduction, often referred to as achieving consistency, is a class of techniques for transforming a CSP into problems that are potentially easier to solve. Problem reduction alone does not normally produce solutions, however it can be extremely useful when used in conjunction with search methods. We define some basic notions of consistency [Tsang, 1993] below.

Node-Consistency of a node : A node is node-consistent if and only if all labels in its domain satisfy all unary constraints on that node.

Arc-Consistency of an arc : An arc (i, j) is arc-consistent if and only if nodes i and j are node-consistent and for any label in the domain of node i, there exists a label in the domain of node j such that all binary constraints on the two nodes are satisfied.

Consistency of a graph : A graph is said to be **node-** or **arc-consistent** if and only if every node or arc respectively is node- or arc-consistent.

Multiple unary constraints on a node can be equivalently represented by a single unary constraint that is their conjunction. The same technique can be applied to represent multiple binary constraints between two nodes as a single binary constraint between them. The graph in Fig. 1 is node consistent as no unary constraints have been specified. However, it is not arc-consistent. We now extend the notion of arc-consistency for non-binary constraints. Instead of having constraints on pairs of nodes defining edges in a constraint graph, we have constraints over n nodes that define the **hyperedges** in a **hypergraph** . We define such constraints as **n-ary** constraints. A constraint is a relation $R_{i,j,\ldots k}$, specifying admissible labels for the nodes $i, j, \ldots k$ over this constraint A labeling is *arc-consistent* if and only if :

$$\forall v \in V, \forall d_v \in D_v, \forall R_{i,j,\ldots k} \text{ constraining } v,$$
$$\forall i, j, \ldots k, \exists d_i, d_j, \ldots d_k \text{ such that}$$
$$R_{i,j,\ldots k}(d_i, d_j, \ldots d_k) \text{ is true}$$

This means that for every label in the domain of every node i, there exist labels in the domains of nodes $j, \ldots k$ such that every *n-ary* constraint on these nodes is satisfied.

3 Optimal Arc-Consistency

We briefly describe the optimal binary arc-consistency algorithm AC4 [Mohr and Henderson, 1986]. Arc-consistency is intuitively based on the notion of *support*. A label at a node is supported by a label at another node if they satisfy the constraint between the two nodes. A label b for a node i is viable if there exists at least one label at each of the other nodes j (connected to it by a constraint) that supports it. In Fig1, label *a* at node 1 is supported by labels *b,c and d* at node 4 over the constraint between nodes 1 and 4. However, *a,b and c* at node 4 do not have any support over the constraint with node 3. Only d at node 4 is supported (by c) over that constraint. If there exists a node j where no label satisfies the required relation with b at node i, then b can be removed as a potential label for node i. In Fig.1, *a* at node 4 is such a label since no labels at node 1 satisfy the constraint with node 4. Fig.2 shows the state of the CSP after achieving arc-consistency. AC4 makes the notion of support explicit by assigning a counter to each arc-label pair. These pairs are denoted

by $[(i,j),b]$ and indicate the arc from i to j with label b at node i. The counters are designated by $counter[(i,j),x]$ and indicate the number of j labels that support the label b for i in the constraint (i,j). In addition, for each label c at node j, the set $S_{jc} = \{(i,b) \mid c$ at node j supports b at node $i\}$ is constructed. If c can no longer be a viable label for node j, then $counter[(i,j),b]$ must be decremented for each b supported by c. A table, M, to keep track of deleted labels and a list $List$ to keep track of constraint propagation are the other structures used. $List$ is initialized with all values (i,b) such that there exists some j for which $counter[(i,j),b] = 0$. The table M ensures that each node-label pair is considered only once. During the propagation phase, the algorithm takes values (j,c) in $List$, decrements each counter $counter[(i,j),b]\ \forall\ (i,b) \in S_{jc}$ and when a $counter[(i,j),b]$ becomes zero, it puts (i,b) in $List$. The algorithm stops when $List$ is empty, which means that all values in M have non-empty supports on all the constraints. The upper bound time complexity for AC4 is $O(ea^2)$ where e is the number of arcs and a is the cardinality of the biggest label set.[Mohr and Henderson, 1986]
GAC4 [Mohr and Masini, 1988] works on the same principle as AC4. A constraint C is a represented as a set of n-tuples $((i,a),(j,b),\ldots(k,c))$. When a label a has to be removed from the set D_i of possible labels for i, all the n-tuples including the label have to be discarded. If the discarded n-tuple is the last in the constraint C supporting a particular label at a node, then that label is deleted and the effect of the deletion is propagated.

4 A New HyperArc-Consistency Algorithm: NAC4

We have seen that constraints in GAC4 are specified as sets of tuples of admissible labels, that is the relation is represented extensionally (explicitly). This may not be feasible in certain problem domains where the constraint relations might be intensionally represented and computing the extension might be very expensive or in some cases unnecessary. Intuitively, the algorithm should take advantage of prior label pruning to minimize the computation of the relation's extension.
It has been shown ([Tsang, 1993] pp.14) that any n-ary constraint can be represented as a set of binary constraints *provided* the constraints are represented as a list of acceptable tuples of labels (as in GAC4). A constraint expressed implicitly such as: node v_1 is "brighter-than" node v_2 or node v_1 is "brighter-than" node v_3 cannot be represented equivalently as a set of binary constraints unless we compute and compare the brightness values of all the labels in the domains of nodes v_1 to v_3.

NAC4 is designed for constraints that are defined by implicit relations. The notion of support, as in AC4, is extended in NAC4. A label at a node is supported by a set of labels, one at each of the other nodes involved in the constraint, and must have a minimum support from these label sets over all the constraints involving it.
We now present the algorithm NAC4 :

Algorithm NAC4

Initialization
1 $M := 0$; $S_{ib} := \emptyset$;
2 **for** $r \equiv (i,j,\ldots k) \in E$ **do**
3 **for** $b \in A_i$ **do**
4 **begin**
5 Total := 0;
6 **for** $(c,\ldots d) \in (A_j \times \ldots \times A_k)$ **do**
7 **if** $R(i,b,j,c,\ldots k,d)$ **then**
8 **begin**
9 Total := Total + 1;
10 $Tuple := \{(j,c),\ldots(k,d)\}$;
11 **for** $(j',c') \in Tuple$ **do**
12 Append$(S_{j'c'},((i,b),r,Tuple))$;
13 **end**;
14 **if** Total = 0 **then**
$M[i,b] := 1$; $A_i := A_i$ - $\{b\}$;
15 **else** $Counter[r,b]$:= Total;
16 **end**;
17 Initialize List with $\{\ (i,b) \mid M[i,b] = 1\ \}$;

Propagation phase
18 **while** List not Empty **do**
19 **begin**
20 choose and remove (j,c) from List;
21 **for** $((i,b),r,T) \in S_{jc}$ **do**
22 **begin**
23 $Counter[r,b] := Counter[r,b]$ - 1;
24 **for** $(k,d) \in T$ **do** $S_{kd} := S_{kd} - ((i,b),r,T)$;
21 **if** $Counter[r,b] = 0$ **and**
$M[i,b] = 0$ **then**
22 **begin**
23 Append(List,(i,b));
24 $M[i,b] := 1$; $A_i := A_i$ - $\{b\}$;
25 **end**;
26 **end**;
27 **end**

In the initialization stage, we construct a list of inconsistent labels, stored in M, by calculating the support for labels at each node over each n-ary relation r. The set $S_{jc} = \{\ (i,b)\}$ was used in AC4 to store the fact that c at j supports b at i. We now augment the information stored in S_{jc} to include the rest of nodes involved in the relation. This is achieved by storing the relation index and a pointer to the whole tuple. $S_{jc} = \{\ ((i,b),r,Tuple) \mid c$ at j supports b at i in the tuple $Tuple$ of the relation $r\ \}$. We therefore append the support list of each node in r whenever it supports label b at i . If a label c were to become incompatible for such a node j, then we

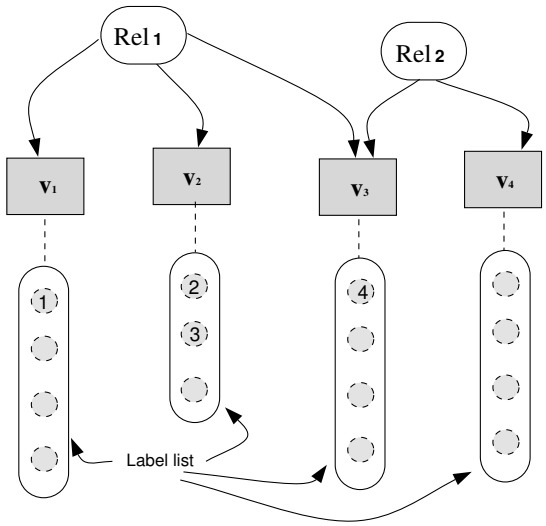

Figure 3: An example of a non-binary CSP

would have a link to re-assess the support for b at i. Counters are indexed over relations but imply support for the label at the $first$ node of the relation.
The propagation phase is altered (line 24) to ensure that "partner" nodes in the same relation r update their support lists. Suppose (j, c) no longer supports (i, b) over relation r, every (k, d) involved in r removes the entry indexed by (i, b) from its support list. Since the counter for b at i over r was incremented only once , multiple decrements are not necessary when $((i, b), r, T)$ is deleted from S_{kd}.

For example, in Fig.3, let us suppose that label 1 at node v_1 is supported by labels 2 and 3 at v_2 and label 4 at v_3 over constraint Rel_1.(In the explicit relation notation, Rel_1 would be $\{(1, 2, 4), (1, 3, 4)\}$) Now suppose label 4 cannot be assigned to v_3 due to some other constraint (possibly Rel_2) and we delete it from the label list of v_3, then we have to update the support of label 1 at v_1. We decrease the counter and delete the corresponding tuple from the support lists of every node-label pair involved in this tuple. Since there is no other label at v_3 supporting the assignment of label 1 to v_1 over Rel_1, label 1 cannot be a valid label for v_1 and must be deleted. The node-label pairs that it supports will now be examined in turn.

5 Analysis

A formal correctness proof is not presented here but can be worked out along the lines of the proof for AC4.

5.1 Correctness of NAC4

We present the steps for a complete proof of the correctness of NAC4.

1. Our induction hypothesis is that each label deleted from A_i is not admissible for any hyperarc-consistency solution. A label is removed if one of its counters goes to zero, so it has no more corresponding labels at one *hyperarc*. By induction, all the previously removed labels could not belong to any solution, so this one cannot belong to any solution.
2. The result of NAC4 is *hyperarc*-consistent: for all hyperarcs r, for all labels b at node i, we have $Counter[r, b] > 0$. Therefore, each b has corresponding label nodes $j, \ldots k$.
3. From the above two steps, we conclude that NAC4 builds the largest hyperarc-consistent solution because only those labels that provably do not participate in solutions to the original problem are removed.

5.2 Complexity

Since the algorithm has been designed for expensive implicit constraints, the domain independent assumption that the constraint verification is of unit cost is invalid. Hence the running time is decided by the application dependent cost of constraint verification. If we were to assume (wrongly) that constraint verification in line 7 of the algorithm is a unit cost operation, the worst case time complexity of the algorithm would be $O(\sum_{n=1}^{m} C_n^m\, na^n)$, [1] where m is the number of nodes and a is the cardinality of the largest domain. However this complexity is not of much practical importance here since this algorithm is designed for constraints whose cost of evaluation is orders of magnitude greater than a unit step of the algorithm. A precise complexity analysis would estimate the relative cost of verifying each constraint relation and weigh the cost summation correspondingly.

6 Expensive Constraints

An expensive constraint is one for which checking whether a tuple belongs to the corresponding relation is computationally expensive. Moreover, this implies that the cost of computing the full extension of a relation is proportionally expensive to the sizes of the domains of the participating nodes. A preprocessing step to order the constraints according to their domain specific costs

[1] For each of the $n \times C_n^m$ n-ary constraint, at most a^n counters are incremented. Since the counters never become negative, at most $(\sum_{n=1}^{m} C_n^m\, na^n)$ decrements are performed.

can be added to any consistency algorithm that represents the constraints implicitly.
For example, the initialization phases of AC4 and NAC4 can be modified to select the next least expensive edge. Such a modification does not seem possible in the case of GAC4 since the initialization uses the explicit representation of the constraints to build the elegant data structures that make the algorithm optimal. The standard way to achieve arc-consistency requires that we first achieve node-consistency. However, some unary constraints may be more expensive than other constraints of higher arity. These unary constraints need not be enforced in scenarios where non-unary constraints are much cheaper to evaluate and achieve substantial problem reduction.

7 Application

The algorithm has been prototyped in Allegro Common Lisp and used for an application in high-level computer vision [Chopra and Srihari, 1995; Srihari *et al.*, 1994]. The nodes are people names, the labels are rectangular sub-images denoting human faces. The constraints are extracted from a caption accompanying the image [Srihari and Burhans, 1994]. Most of the constraints are binary spatial and visual relations but some are also n-ary spatial relations. Another set of constraints include unary constraints that are orders of magnitude more expensive to compute than the spatial constraints. Traditional use of CSP techniques in computer vision assume that features are first extracted and represented as nodes of the CSP with semantic labels being the values. Since we have a strong hypothesis, from the caption, of the image contents in our application, we reverse the roles of features and semantic labels in the CSP (as indicated in the beginning of this paragraph). Therefore, features are extracted and constraint relations are verified only when required.
No extensive experimentation on randomly generated data has been conducted using this method for n-ary constraints. The focus has been on providing a version of arc-consistency capable of handling n-ary constraints specified through implicit relations.

8 Conclusion and Future Work

We have provided an algorithm that handles certain types of n-ary constraints and that doesn't make the assumptions of GAC4 regarding representation of the constraints. We believe the assumption of unit time complexity assigned to checking a constraint is invalid and in some cases the constraint check time easily outweighs the time spent in the rest of the algorithm. Hence the complexity of the algorithm is dependent upon the problem domain that defines the constraint checking time. Domain specific experiments or simulations need to be performed to validate our belief in the better performance of NAC4 with implicit relations over GAC4 with explicit relations. A rigorous complexity analysis with reasonable estimates of the relative costs of checking different constraints needs to be performed.

9 Acknowledgments

This work was supported in part by ARPA Contract 93-F148900-000. We would like to thank Kannan Govindarajan and Hans Chalupsky for comments and help in writing this paper.

References

[Bessiere, 1991] C. Bessiere. Arc consistency in dynamic constraint satisfaction problems. In *Proceedings of AAAI-91*, pages 221–226, 1991.

[Chopra and Srihari, 1995] Rajiv Chopra and Rohini Srihari. Control structures for incorporating picture-specific context in image interpretation. In *Proceedings of IJCAI-95*, Montreal,Canada, August 1995. Morgan Kaufmann.

[Faugeras and Berthod, 1981] O. Faugeras and M. Berthod. Improving consistency and reducing ambiguity in stochastic labelling: an optimization approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4:412–423, 1981.

[Haralick and Elliot, 1980] R.M. Haralick and G.L. Elliot. Increasing tree search efficiency for constraint satisfaction problems. *Artificial Intelligence*, 14:263–313, 1980.

[Hentenryck *et al.*, 1992]
P. Van Hentenryck, Y. Deville, and C. Teng. A generic arc-consistency algorithm and its specializations. *Artificial Intelligence*, 57:291–321, 1992.

[Mackworth and Freuder, 1985] A.K. Mackworth and E.C. Freuder. The complexity of some polynomial network consistency algorithms for the constraint satisfaction problem. *Artificial Intelligence*, 25:65–74, 1985.

[Mackworth, 1977] A.K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.

[Mohr and Henderson, 1986]
R. Mohr and Thomas C. Henderson. Arc and path consistency revisited. *Artificial Intelligence*, 28:225–233, 1986.

[Mohr and Masini, 1988] R. Mohr and G. Masini. Good old discrete relaxation. In *Proceedings of ECAI-88*, pages 651–656, 1988.

[Srihari and Burhans, 1994] Rohini K. Srihari and Debra T. Burhans. Visual Semantics: Extracting Visual Information from Text Accompanying Pictures. In *Proceedings of AAAI-94*, pages 793–798, 1994. Seattle, WA.

[Srihari *et al.*, 1994] Rohini Srihari, Rajiv Chopra, Debra Burhans, Venugopal Govindaraju, and Mahesh Venkatraman. Use of collateral text in image interpretation. In *Proc. of the ARPA Image Understanding Workshop*, pages 897–908, November 1994.

[Tsang, 1993] Edward Tsang. *Foundations of Constraint Satisfaction*. Academic Press, 1993.

[Waltz, 1975] D. Waltz. Understanding line drawings of scenes with shadows. In P.H. Winston, editor, *The Psychology of Computer Vision*. McGraw-Hill, New York, 1975.