Using Interval Arithmetic To Model Finite Domain CSPs Where Domain Generation Is Expensive

Rajiv Chopra and Rohini K. Srihari
Department of Computer Science and Center for Document Analysis and Recognition
State University of New York at Buffalo, Buffalo, New York 14228-2567, U. S. A.
{rchopra, rohini}@cedar.buffalo.edu

Abstract

It is generally assumed that the variables, domains and constraints of a Finite-Domain Constraint Satisfaction Problem are all pre-computed inputs to a *black-box* constraint satisfaction algorithm. The obvious advantages of such an assumption is the freedom in developing generic constraint solvers and the declarative use of constraint technology. However, it is useful to examine applications of CSPs to evaluate if they can benefit further from *on-the-fly* constraint and domain specification. We consider particular classes of numerically computable discrete CSPs for which the variables' domain generation is costly. The domain generation and its cost are usually application dependent. However, in several numeric discrete CSPs, some general measures of cost can be estimated. We propose to reduce the generation cost of the discrete finite domains by combining interval narrowing techniques with local consistency techniques in discrete domains. We present a new algorithm HetCSP for the above task and illustrate its use with an example and several experimental results.

Introduction and Definitions

In traditional Constraint Satisfaction Problems, it is generally assumed that the variables, domains and constraints are all pre-computed inputs to a *black-box* constraint satisfaction algorithm. The obvious advantages of such an assumption is the freedom in developing generic constraint solvers and the declarative use of constraint technology. However, it is useful to examine applications of CSPs to evaluate if they can benefit further from *on-the-fly* constraint and domain specification. We consider particular classes of numerically computable discrete CSPs for which the variables' domain generation is costly. The domain generation and its cost are usually application dependent. However, in several numeric discrete CSPs, some general measures of cost can be estimated. We propose to reduce the generation cost of the discrete finite domains by combining interval narrowing techniques with local consistency techniques in discrete domains.

The motivation for this work comes from an application in high-level computer vision(Chopra & Srihari 1995; Srihari *et al.* 1996) where the constraint graph consisted of a hypothesis of the contents of an image. The variables of the CSP are physical objects, for instance, buildings or faces, and spatial relations between these objects formed the constraints. They presented a iterative technique called *spatial prediction* that demarcated subregions of the original image as *search regions*. Computer vision algorithms would then be applied selectively on these sub-images to generate potential values for the variables. However, there were several problems with their approach as described later in the Discussion section. We believe our algorithm overcomes those problems and can easily be applied to this application as well as others.

Definition 1: A Discrete Constraint Satisfaction Problem (DCSP), (V, D, C, R), involves a set of n **variables**, $V = \{v_1, v_2, \ldots v_n\}$, each with an associated **domain** D_i of possible **values**. The search space consists of the Cartesian product of the variables' domains, $D = D_1 \times D_2 \times \ldots \times D_n$. A set of **constraints** C is specified over some subsets of these variables, where each constraint C_p involves a subset $\{i_1, \ldots, i_q\}$ of V and is labeled by a relation R_p of R, subset of the Cartesian product $D_{i_1} \times \ldots \times D_{i_q}$ that specifies which values of the variables are compatible with each other. A **solution** to a DCSP consists of an assignment of values to variables such that all constraints are satisfied.

Various approaches exist for solving discrete CSPs. A common approach is based on problem reduction strategies, and involves achieving local, that is, node or arc, consistency (Mackworth 1977; Mohr & Henderson 1986; Mohr & Masini 1988; Hentenryck, Deville, & Teng 1992). This reduces the amount of backtracking required for global consistency.

Definition 2: An interval is informally defined as a number pair: $X = [a : b] = \{x | a \leq x \leq b\}$

Definition 3: We define a set G of functions g_i called domain generators. The function g_i associated with variable v_i maps an interval to a discrete finite domain for that variable or equivalently $g_i(I_i) \rightarrow D_i$.

These functions are typically application dependent and we focus on those applications where these generators are computationally expensive. We present below one measure of the cost associated with generating the domains. For example, in the vision application mentioned in the introduction, the object location tools would be the generators.

Definition 4: Since our objective is to reduce the cost of generating the domains of the variables, let p_i be the cost associated with generating the domain of variable v_i. We define *cost* to be a function of g_i and I_i implying that $p_i = cost(g_i, I_i)$. The total cost of generating the domains is $\mathcal{P} = \sum_{i=1}^{n} p_i$ and our objective is $min(\mathcal{P})$. If we assume that *cost* is directly proportional to $|I|$, then $max(p_i) = cost(g_i, max(|I_i|))$. In solving the CSP, assuming that a solution exists, all the domains will be generated. Note that in a traditional discrete finite domain CSP, all domains are generated initially with intervals being their maximum respective values, leading to a cost $\mathcal{P} = \sum max(p_i)$. Any algorithm that reduces the size of the intervals (assuming that the generator functions cannot be changed) achieves a reduction in $\mathcal{P}$.

Definition 5: We define our new heterogeneous CSP (V, G, C, P, V', D', C') in the following manner: The discrete finite-domain constituent (FDCSP) is a triple (V, G, C) where V is a set of n variables $\{v_1, v_2, \ldots, v_n\}$, G is a set of generators g_i that generate domains D_i for the n variables and C is a set of constraints where C_p involves a subset $\{i_1, \ldots, i_q\}$ of V and is labeled by an *implicit* relation R_p of R, subset of that specifies which values of the variables are compatible with each other.

The other constituent of our heterogeneous CSP is an interval CSP (ICSP) defined by the triple (V', D', C'). For each variable $v \in V$ there exists a variable $v' \in V'$. D' is a set of interval domains where D'_i denotes the domain for the interval variable v'_i. For example, if D_i, the domain of the FDCSP variable v_i is a set of reals, then the domain D'_i of the corresponding interval variable v'_i is an interval over reals, I_i as described by Definition 2. The set P comprises of the domain generators p_i defined above in Definition 3. The set of constraints C' consists of implicit relations corresponding to the relations over the FD variables. The semantics of these relations is the arena of interval arithmetic and is outlined in Appendix A.

Definition 6: We define $\mathbf{hull}(S) = [min(S) : max(S)]$ where S is a finite set of values from a totally ordered set.

Algorithm HetCSP

We now present the sketch of an algorithm to reduce the cost of generating the domains of the FDCSP in an incremental manner.

Algorithm HetCSP

A0: Create an incremental FDCSP ($V'' = NULL$, $D'' = NULL, C'' = NULL$)
A1: Ensure local consistency in ICSP (V', D', C')
A2: Choose a variable $v_i \in V - V''$ such that $p_i = min[p(g_i, I_i)]$
A3: Generate the finite domain D''_i and update the FDCSP such that $V'' = V'' \cup \{v_i\}$, $D'' = D'' \cup \{D''_i\}$ and $C'' = C'' \cup \bigcup_{v_j \in V''} \{c'' | c'' = \rho(x_1, \ldots, x_i, \ldots, x_k) \wedge x_1, \ldots, x_k \in V''\}$
A4: Ensure local consistency in the FDCSP
A5: $\forall v_i \in V''$ compute domains D'_i of corresponding variables v'_i such that $D'_i = \mathbf{hull}(D''_i)$
A6: If $V - V'' = NULL$ quit, else **goto A1**.

Local consistency in the ICSP is ensured by one of several interval narrowing techniques(Siskind & McAllester 1993b; Benhamou, McAllester, & Hentenryck 1994; Benhamou & Older 1993; Faltings 1994; Sam-Haroud & Faltings 1996). Local consistency algorithms in the FDCSP is ensured to different degrees by arc- and path- consistency algorithms(Mackworth 1977; Mohr & Henderson 1986; Hentenryck, Deville, & Teng 1992). This algorithm falls under the class of consistency or problem reduction algorithms and requires a final enumeration procedure (e.g. chronological backtracking) for completeness.

Consider the example CSP shown in Figure 1. Frame 1 describes the initial state of the CSP with 3 finite domain variables A, B and C and their corresponding interval variables A', B' and C'. The initial interval from which the discrete finite values for A, B and C are chosen is [1:100]. The empty curly braces denote the finite domains.

Frame 2 shows the state of the CSP once the domain of A is generated. The interval is accordingly recomputed (Step A5) to [20:30] and this effect is propagated over the interval variables by narrowing B' and C' to [20:100] (Step A1).

In frame 3, the domain of variable C was generated and set to {22,60,80}. The interval C' is according narrowed to [22:80] and in the propagation step, B' is narrowed to [20:80].

Finally, in frame 4, the domain of B is generated. It is set to {25}. Arc-consistency applied to the FDCSP removes 30 from the domain of A and 22 from the domain of C. The interval domains are narrowed to reflect the final state. Since only local consistency was ensured in the incremental stages, a final pass to

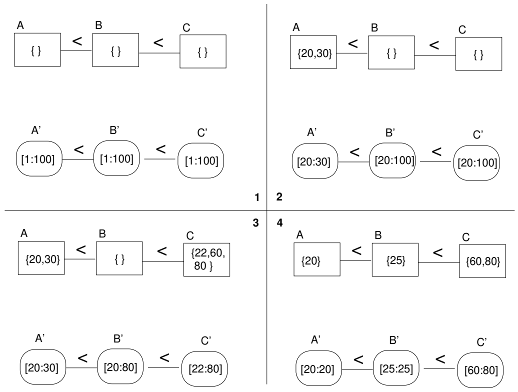

Figure 1: Example: Frames 1 to 4 outline the state of the CSP with 3 finite domain variables A, B and C and their corresponding interval variables A', B' and C'

achieve global consistency has to be performed and solutions enumerated.

Analysis

It is easy to see that the variable selection in Step A2 makes HetCSP a greedy algorithm. Several different heuristics could be substituted here and need to be investigated further. The focus of this algorithm is on invoking the domain generators with the minimum cost. We chose the greedy approach as it is the simplest and we make no claims about optimality.

Correctness

The correctness of HetCSP is informally proven as follows:

1. Routines for local consistency in ICSP and FDCSP exist and are correct. In other words, the local consistency routines always terminate and filter the domains (interval and discrete-finite) such that no value that participates in a correct solution is removed.

2. Discrete values that are deleted from the domains of variables are done so only by the (correct) local consistency algorithm for FDCSPs.

3. Discrete values for a variable that are not generated as a result of the corresponding interval being narrowed cannot potentially participate in a correct solution. They would be in conflict with generated values at other nodes that caused this particular interval to be narrowed.

4. HetCSP terminates trivially since it iterates n times over steps (A1-A6) that have only terminating routines.

5. HetCSP is correct since it terminates and deletes only those values from the domains of the FD variables that cannot satisfy the constraints and participate in a correct solution.

Implementation and Experimentation

This algorithm has been implemented in Allegro Common Lisp and uses a constraint package included with the public domain software *SCREAMER* (Siskind & McAllester 1993b), for enforcing local consistency on interval variables.

The algorithm was tested on randomly generated CSPs. The generation of the finite domains was simulated. The routine to generate random CSPs took four parameters: number of variables in the CSP : *num-var*, mean number of values for each variable : *num-val*, continuous interval from which discrete values are chosen: *int-domain*, and approximate variance in the number of values : *var-val*

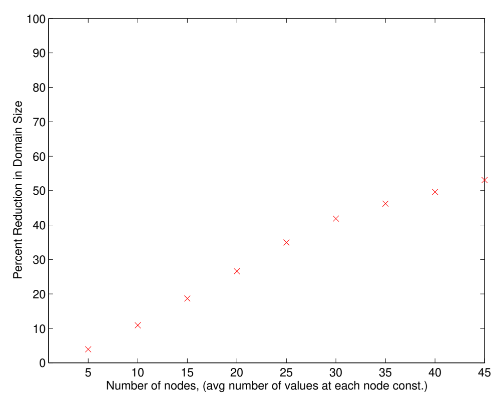

Figure 2: Experiment 1: Each point represents 100 randomly generated CSPs that had at-least one or more possible labelings. The mean number of values at each node was 30 and the approximate variance was 3. Since the mean number of values at each node is constant and the number of constraints to be satisfied increases, there is an increase in the reduction of the domain sizes.

The output of the above routine is a triplet of arrays:
IntDomains : an array of num-var interval variables initialized to the interval *int-domain*
FDDomains : an array of num-var discrete finite sets of values within the interval. This is the simulated *Generator* of finite domains.
Constraints : a 2D array (size num-var $\times$ num-var) of constraints between variables. For simplicity, in this

experiment we define a single constraint $<$. Without any loss in generality, we generate the constraints such that if Constraints[i,j] exists, then $i < j$. Since the $<$ relation is transitive, this restriction prevents cycles in the constraint relationships. Note however that $i < j$ does not imply that a constraint exists between them. The assignment of constraints is randomly generated.

Figures 2 to 5 display the results of four experiments that were performed to analyze the algorithm. We present the percentage reduction in domain size versus several parameters. Each point represents the average of 100 randomly generated CSPs that had at-least one or more possible labelings. For each CSP, the original domain size is calculated as the sum of the sizes of each variable's domain. If the order of magnitude of the cost of generating the domain of a variable is $\Theta(\mid I_i \mid)$ then the percentage reduction in domain size gives the reduction in generation cost. If the generation cost is proportional to higher powers of the domain size, then the savings are greater. In all the experiments the intervals are initialized to [0:100].

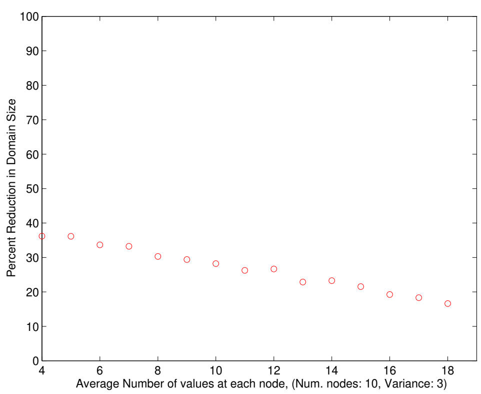

Figure 3: Experiment 2: Each point represents 100 randomly generated CSPs that had at-least one or more possible labelings. Since the interval size is constant, as the mean number of values per node increases, the reduction in the domain size decreases.

In Exp. 1, the mean number of values at each node was 30 and the approximate variance was 3. Since the mean number of values at each node is constant and the number of constraints to be satisfied increases, there is an increase in the reduction of the domain sizes. In experiment 2, as the average number of values per node increases, the reduction in the domain sizes decreases. Since the initial interval is constant, a larger number of values from the same interval leads to larger intermediate intervals and hence a lower total reduction in domain sizes.

We measured the effect of increasing the number of variables in experiment 1 . However, there is a side

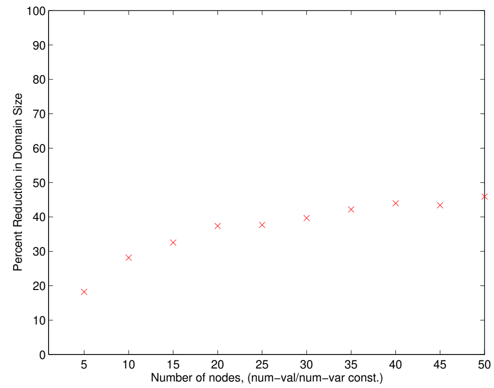

Figure 4: Experiment 3: Each point represents 100 randomly generated CSPs that had at-least one or more possible labelings. By keeping the mean number of values per variable constant, we are able to counter the effect of fewer values per variable and measure how the algorithm scales up with increasing number of variables (in contrast with Experiment 1).

effect of the average number of values per variable decreasing. Therefore, in Experiment 3, to counter balance that side effect, we increase the number of values per variables as we increase the number of variables. We find that the reduction in domain sizes increases at a slower rate than in Exp. 1. We are also able to measure the scalability of the algorithm with increasing number of variables. We find through Experiment 4 that the variability in the relative sizes of the domains does not affect the domain reduction. Since the values are randomly chosen from the same initial interval domain, fewer values does not imply smaller interval sizes.

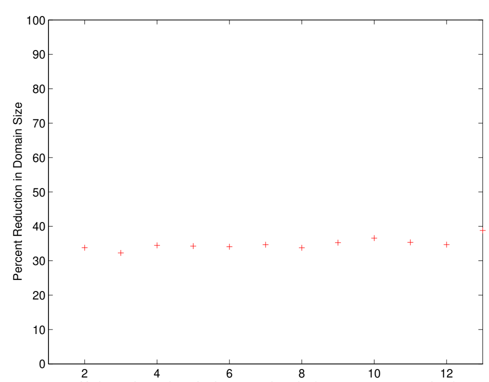

Figure 5: Experiment 4: Each point represents 100 randomly generated CSPs that had at-least one or more possible labelings. The variability in the relative sizes of the domains does not affect the domain reduction.

Discussion

A different viewpoint on domain generation is to consider an discrete constraint satisfaction problem where the domains of the variables are extremely large and the domain generators are unary constraints on these variables. However, the cost of evaluating the unary constraint is directly proportional to some positive power of the size of the domain. This view treats the problem as a discrete CSP only, whereas we view it as a combination of interval and discrete CSPs, where interval narrowing on numerical domains is computationally less expensive.

A potential problem in applying this algorithm may lie in the estimation of the domain generation cost. We make the assumption that the discrete numeric values that constitute the finite domains of the variables are generated from some numeric interval (as defined in Definition 3). We further assume that the generation cost is application specific, though proportional to some positive power of the size of the interval. It is this cost that we attempt to minimize. The observation that a single interval narrowing step can drastically reduce the interval representing the finite domain is precisely what helps reduce the (costlier application-specific) domain generation cost. Interval narrowing techniques have been used in association with finite domain CSPs in a layout generation application (Hower 1996). The objective there was to compute all layout possibilities in a semi-automatic system where a user places non-overlapping objects in a two dimensional frame. The application is an excellent example of the benefits of reformulating discrete finite domain CSPs using interval constraints. However, domain generation was not an issue, unlike this paper. Can partial results be used to guide search? In the past, backtrack search algorithms have incorporated domain independent techniques and heuristics (Chapters 5-7 (Tsang 1993)) to guide the search for solutions using partial results. However, all the above techniques are for CSPs where the domains and constraints are pre-computed or can be computed with unit cost. We believe our algorithm is the first to consider domain generation costs as minimization criteria for CSPs.

Chopra and Srihari (Chopra & Srihari 1995; Srihari *et al.* 1996) explored domain generation in a computer vision application where the constraint graph consisted of a hypothesis of the contents of an image. For example, the variables were buildings or faces and spatial relations between these objects formed the constraints. They presented a technique called *spatial prediction* that demarcated subregions of the original image as *search regions*. Computer vision algorithms would then be applied selectively on these sub-images to generate potential values for the variables. The problems with the above mentioned approach are:
1) Incomplete propagation: When *spatial prediction* does take place, it changes the search regions of **only** the variables adjacent to it in the constraint graph. The change in the search regions is not propagated to other variables' search regions.
2) Limited activation of propagation: Spatial prediction takes place only when a variable is uniquely labeled. If multiple values are generated for a variable, and partial satisfaction of the constraint graph does not lead to unique labeling for any variable, no spatial prediction takes place.
3) Domain dependence of algorithm: The above algorithm uses data structures and spatial terms particular to the high level vision problem addressed.

Algorithm HetCSP addresses the three weaknesses described above. Propagation of the interval bounds takes place in all interval variables till local consistency is achieved in the ICSP. Step A5 computes new interval bounds and propagation is activated every time a domain changes; not just when there is a single value in the domain. The presented algorithm is independent of the application domain and a mapping from the vision application to HetCSP is conceivable though not presented here.

Conclusion

We have presented an innovative application of interval arithmetic in constraint satisfaction. Though interval constraint satisfaction is an extensive research area, we believe this to be the first attempt at using interval arithmetic to reduce the domain generation cost of finite domain CSPs. A new algorithm has been proposed here for the above. Illustration with a simple example, analysis, implementation and experimentation of the algorithm have also been detailed here. We believe the algorithm will prove useful in the application domain that motivated this research and also believe that the presentation here is general enough for application in other areas.

References

Alefeld, G., and Herzberger, J. 1983. *Introduction to Interval Computations.* Academic Press.

Benhamou, F., and Older, W. J. 1993. Applying Interval Arithmetic to Real, Integer and Boolean Constraints. *Logic Programming: The ALP Newsletter* 6(2):13–14. (Extended Abstract).

Benhamou, F.; McAllester, D.; and Hentenryck, P. V. 1994. CLP(Intervals) Revisited. In *Logic Programming - Proceedings of the 1994 International Symposium*, 124–138.

Chopra, R., and Srihari, R. 1995. Control Structures for Incorporating Picture-Specific Context in Image Interpretation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI95)*, 50–55. Montreal,Canada: Morgan Kaufmann.

Davis, E. 1987. Constraint Propogation with Interval Labels. *Artificial Intelligence* 32(3):281–331.

Faltings, B. 1994. Arc-consistency for continuous variables. *Artificial Intelligence* 65(2):363–376.

Hentenryck, P. V.; Deville, Y.; and Teng, C. 1992. A Generic Arc-Consistency Algorithm and its Specializations. *Artificial Intelligence* 57:291–321.

Hower, W. 1996. Bottom-up Layout Generation. *Informatica* 20(1):57–63.

Hyvönen, E. 1989. Constraint Reasoning Based on Interval Arithmetic. In *Proceedings of IJCAI 1989*, 1193–1198.

Hyvönen, E. 1992. Constraint Reasoning Based on Interval Arithmetic: The Tolerance Propogation Approach. *Artificial Intelligence* 58:71–112.

Mackworth, A. 1977. Consistency in Networks of Relations. *Artificial Intelligence* 8:99–118.

Mohr, R., and Henderson, T. C. 1986. Arc and Path Consistency Revisited. *Artificial Intelligence* 28:225–233.

Mohr, R., and Masini, G. 1988. Good Old Discrete Relaxation. In *Proceedings of ECAI-88*, 651–656.

Moore, R. E. 1966. *Interval Analysis.* Prentice Hall.

Sam-Haroud, D., and Faltings, B. 1996. Consistency Techniques for Continuous Constraints. *Constraints Journal* 1.

Siskind, J. M., and McAllester, D. A. 1993a. Nondeterministic Lisp as a Substrate for Constraint Logic Programming. In *Proceedings of the Eleventh National Conference on Artifical Intelligence.*

Siskind, J. M., and McAllester, D. A. 1993b. Screamer: A Portable Efficient Implementation of Nondeterministic Commonlisp. Technical Report IRCS–93–03, University of Pennsylvania Institute for Research in Cognitive Science.

Srihari, R.; Zhang, Z.; Venkatraman, M.; and Chopra, R. 1996. Using Speech Input for Image Interpretation and Annotation. In *Proceedings of the ARPA Image Understanding Workshop.*

Tsang, E. 1993. *Foundations of Constraint Satisfaction.* Academic Press.

Interval Arithmetic and Constraints

We present the basic notions of rational interval arithmetic (IA) (Moore 1966), and interval arithmetic applied to constraint satisfaction (Davis 1987; Benhamou & Older 1993; Hyvönen 1989; 1992).

We consider $\Re \bigcup \{-\infty, +\infty\}$, the set of real numbers augmented with the two infinity symbols, and the natural extension of the relation $\leq$ to this set. The **interval** is defined as a set of real numbers represented as a number pair. For every $a, b \in \Re, a \leq b$, we use the following notations for intervals:
$[a:b] = \{x \in \Re | a \leq x \leq b\}, [a:b) = \{x \in \Re | a \leq x < b\}$
$(a:b] = \{x \in \Re | a < x \leq b\}, (a:b) = \{x \in \Re | a < x < b\}$
$[a:+\infty) = \{x \in \Re | a \leq x\}, (a:+\infty) = \{x \in \Re | a < x\}$
$(-\infty:b] = \{x \in \Re | x \leq b\}, (-\infty:b) = \{x \in \Re | x < b\}$
$(-\infty:+\infty) = \Re$
In IA, basic algebraic relations can be defined as:
$A + B = \{a + b | a \in A, b \in B\}$
$A - B = \{a - b | a \in A, b \in B\}$
$A * B = \{a * b | a \in A, b \in B\}$
$A/B = \{a/b | a \in A, b \in B\}, if 0 \notin B$
For the above set of relations we can define the limits of the resultant interval in terms of the original intervals as :
$[a:b] + [c:d] = [a+c:b+d]$
$[a:b] - [c:d] = [a-d:b-c]$
$[a:b] * [c:d] = [min(ac, ad, bc, bd) : max(ac, ad, bc, bd)]$
$[a:b]/[c:d] = [a:b] * [1/d:1/c], if 0 \notin [c:d]$
For a detailed description of the properties of interval arithmetic, one can refer (Moore 1966; Alefeld & Herzberger 1983). The definition of *solution functions* or *narrowing functions* (Benhamou & Older 1993) associated with every n-ary relation of $\Re$ is the basis of interval constraint satisfaction. The main properties of the narrowing functions are i) *Contractance:* The narrowed intervals are smaller than the initial intervals, ii) *Correctness:* Every valid solution in the theoretical real numbers lies in the narrowed intervals, iii) *Monotonicity:* The narrowing preserves the inclusion and iv) *Idempotence:* The narrowed intervals have to be computed only once. Several algorithms based on the iterative application of narrowing functions have been proposed in the literature (Davis 1987; Benhamou & Older 1993; Hyvönen 1989; Siskind & McAllester 1993a; Benhamou, McAllester, & Hentenryck 1994; Faltings 1994; Sam-Haroud & Faltings 1996). The differences in the algorithms lies primarily in the set of relations over $\Re$ for which practical narrowing functions have been constructed and in the representation of the space of solutions. Invariably the algorithms terminate trivially and are independent of the order of application of the constraints.