

# Nondeterministic Lisp as a Substrate for Constraint Logic Programming

Jeffrey Mark Siskind\*

University of Pennsylvania IRCS  
3401 Walnut Street Room 407C  
Philadelphia PA 19104  
215/898-0367  
internet: Qobi@CIS.UPenn.EDU

David Allen McAllester†

M. I. T. Artificial Intelligence Laboratory  
545 Technology Square Room NE43-412  
Cambridge MA 02139  
617/253-6599  
internet: dam@AI.MIT.EDU

## Abstract

We have implemented a comprehensive constraint-based programming language as an extension to COMMON LISP. This constraint package provides a unified framework for solving both numeric and non-numeric systems of constraints using a combination of local propagation techniques including binding propagation, Boolean constraint propagation, generalized forward checking, propagation of bounds, and unification. The backtracking facility of the nondeterministic dialect of COMMON LISP used to implement this constraint package acts as a general fallback constraint solving method mitigating the incompleteness of local propagation.

## Introduction

Recent years have seen significant interest in constraint logic programming languages. Numerous implementations of such languages have been described in the literature, notably CLP( $\mathcal{R}$ ) (Jaffar and Lassez 1987) and CHIP (Van Hentenryck 1989). The point of departure leading to these systems is the observation that the unification operation at the core of logic programming can be viewed as a method for solving equational constraints between logic variables which range over the universe of Herbrand terms. A natural extension of such a view is to allow variables to range over other domains and augment the programming language to support the formulation and solution of systems of constraints appropriate to these new domains. The

notion of extending a programming language to support constraint-based programming need not be unique to logic programming. In this paper we present the constraint package included with SCREAMER, a nondeterministic dialect of COMMON LISP described by Siskind and McAllester (1993). This package provides functionality analogous to CLP( $\mathcal{R}$ ) and CHIP in a COMMON LISP framework instead of a PROLOG one.

SCREAMER augments COMMON LISP with the capacity for writing nondeterministic functions and expressions. Nondeterministic functions and expressions can return multiple values upon backtracking initiated by failure. SCREAMER also provides the ability to perform *local* side effects, ones which are undone upon backtracking. Nondeterminism and local side effects form the substrate on top of which the SCREAMER constraint package is constructed.

## Variables and Constraints

SCREAMER includes the function `make-variable` which returns a data structure called a *variable*. SCREAMER variables are a generalization of PROLOG logic variables. Initially, new variables are unbound and unconstrained. Variables may be bound to values by the process of solving constraints asserted between sets of variables. Both the assertion of constraints and the ensuing binding of variables is done with local side effects. Thus constraints are removed and variables become unbound again upon backtracking.

SCREAMER provides a variety of primitives for constraining variables. Each constraint primitive is a “constraint version” of a corresponding COMMON LISP primitive. For example, the constraint primitive `+v` is a constraint version of `+`. An expression of the form `(+v x y)` returns a new variable  $z$ , which it constrains to be the sum of  $x$  and  $y$  by adding the constraint  $z = (+ x y)$ . By convention, a SCREAMER primitive ending in the letter `v` is a constraint version of a corresponding COMMON LISP primitive. Table 1 lists the constraint primitives provided by SCREAMER. All of these primitives have the property that they accept variables as arguments—in addition to ground

\*Supported in part by a Presidential Young Investigator Award to Professor Robert C. Berwick under National Science Foundation Grant DCR-85552543, by a grant from the Siemens Corporation, by the Kapor Family Foundation, by ARO grant DAAL 03-89-C-0031, by DARPA grant N00014-90-J-1863, by NSF grant IRI 90-16592, and by Ben Franklin grant 91S.3078C-1

†Supported in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-91-J-4038.

Type Restrictions:	numberpv realpv integerpv booleanpv memberv
Boolean:	andv orv notv
Numeric:	<v <=v >v >=v =v /=v +v -v *v /v minv maxv
Expression:	equalv
Functions:	funcallv applyv

Table 1: The constraint primitives provided by SCREAMER.

values—and return a variable as their result. The constraint primitive installs a constraint between the arguments and the returned variable stating that, under any interpretation of the variables involved, the value of the result variable must equal the corresponding COMMON LISP primitive applied to the values of the arguments. As another example, the expression (`<v x y`) returns a variable  $z$  and adds the constraint  $z = (< x y)$ . This constraint is satisfied when  $z$  is either `t` or `nil` depending on whether  $x$  is less than  $y$ . For the most part, each constraint primitive obeys the same calling convention as the corresponding COMMON LISP primitive. SCREAMER performs a variety of optimizations to improve run time efficiency. In particular, if the value of a variable returned by a constraint primitive can be determined at the time the function is called then that value is returned directly without creating a new variable.

In SCREAMER, most constraints are of the form  $z = (f x_1 \dots x_n)$  where  $f$  is the COMMON LISP primitive corresponding to some constraint primitive. Constraints of this form can imply type restrictions on the variables involved. For example, a constraint of the form  $z = (< x y)$  implies that  $z$  is “Boolean”, i.e., either `t` or `nil`. Furthermore, this constraint implies that  $x$  and  $y$  are numeric. In practice, a variable usually has a well defined type, e.g., it is known to be Boolean, known to be real, known to be a cons cell, etc. Knowledge about the type of a variable has significant ramifications for the efficiency of SCREAMER’s constraint satisfaction algorithms. SCREAMER has special procedures for inferring the types of variables. Because knowledge of the types of variables is important for efficiency, in contrast to the COMMON LISP primitives `and`, `or`, and `not` which accept any arguments of any type, the SCREAMER constraint primitives `andv`, `orv`, and `notv` require their arguments to be Boolean. This allows the use of Boolean constraint satisfaction techniques for any constraints introduced by these primitives. Similarly, constraint “predicates” return Boolean variables. For example, in contrast to the COMMON LISP primitive `member` which can return the sub-list of the second argument whose head satisfies the equality check, the result of the `memberv` primitive is constrained to be Boolean.

SCREAMER includes the primitive `assert!` which

can be used to add constraints other than those added by the constraint primitives. Evaluating the expression (`assert! x`) constrains  $x$  to equal `t`. This can be used in conjunction with other constraint primitives to install a wide variety of constraints. For example, (`assert! (<v x y)`) effectively installs the constraint that  $x$  must be less than  $y$ .<sup>1</sup> Certain constraint primitives in table 1, in conjunction with `assert!`, can be used to directly constrain the type of a variable. For example, evaluating (`assert! (numberpv x)`) effectively installs the constraint that  $x$  must be a number. Likewise evaluating (`assert! (booleanpv x)`) installs the constraint that  $x$  must be Boolean. This is effectively the same as evaluating (`assert! (memberv x '(t nil))`).

All constraints in SCREAMER are installed either by `assert!` or by one of the constraint primitives in table 1. A constraint installed by `assert!` states that a certain variable must have the value `t`. A constraint installed by a constraint primitive always has the form  $z = (f x_1 \dots x_n)$  where  $f$  is either a COMMON LISP primitive or a slight variation on a COMMON LISP primitive. The variations arise for constraint primitives such as `orv` and `memberv` where the semantics of the constraint version differs slightly from the semantics of the corresponding COMMON LISP primitive as discussed above.

An attempt to add a constraint fails if SCREAMER determines that the resulting set of constraints would be unsatisfiable. For example, after evaluating (`assert! (<v x 0)`) a subsequent evaluation of (`assert! (>v x 0)`) will fail. A call to a constraint primitive can fail when it would generate a constraint inconsistent with known type information. For example, if  $x$  is known to be Boolean then an evaluation of (`+v x y`) will fail.

## Constraint Propagation

In this section we discuss the five kinds of constraint propagation inference processes performed by SCREAMER. First, SCREAMER implements binding propagation, an incomplete inference technique sometimes called value propagation. Second, SCREAMER implements Boolean constraint propagation (BCP). This is an incomplete form of Boolean inference that can be viewed as a form of unit resolution. Third, SCREAMER implements generalized forward checking (GFC). This is a constraint propagation technique for discrete constraints used in the CHIP system. Fourth, SCREAMER implements bounds propagation on numeric variables. Such bounds propagation—when combined with the divide-and-conquer technique

<sup>1</sup>To mitigate the apparent inefficiency of this conceptually clean language design, the implementation optimizes most calls to `assert!`, such as the calls (`assert! (notv (realpv x))`) and (`assert! (<=v x y)`), to eliminate the creation of the intermediate Boolean variable(s) and the resulting local propagation.

discussed later in this paper—implements a generalization of the interval method of solving systems of nonlinear equations proposed by Hansen (1968). Finally, SCREAMER implements unification. Unification is viewed as a constraint propagation inference technique which can be applied to equational constraints involving variables that range over S-expressions. The constraint propagation techniques are incrementally run to completion whenever a new constraint is installed by `assert!` or one of the constraint primitives. The five forms of constraint propagation are described in more detail below.

Each form of constraint propagation can be viewed as an inference process which locally derives information about variables. All forms of propagation are capable of inferring values for variables. For example, after evaluating (`assert! (orv x y)`) and (`assert! (notv x)`) BCP will infer that  $y$  must have the value `t`. If some constraint propagation inference process has determined a value for some variable  $x$  then we say that  $x$  is *bound* and the inferred value of  $x$  is called the *binding* of  $x$ .

**Binding Propagation:** As noted above, most constraints in SCREAMER are of the form  $z = (f x_1 \dots x_n)$  where  $f$  is a COMMON LISP primitive,  $z$  is a variable, and each  $x_i$  is either a variable or a specific value. For any such constraint SCREAMER implements a certain value propagation process. More specifically, if bindings have been determined for all but one of the variables in the constraint, and a binding for the remaining variable follows from the constraint and the existing bindings, then this additional binding is inferred. This general principle is called *binding propagation*. Binding propagation will always bind the output variable of a constraint primitive whenever the input variables become bound. For example, given the constraint  $z = (+ x y)$ , if  $x$  is bound to 2 and  $y$  is bound to 3, then binding propagation will bind  $z$  to 5. Often, however, binding propagation will derive a binding for an input from a binding for the output. For example, given the constraint  $z = (+ x y)$ , if  $z$  is bound to 5 and  $x$  is bound to 2, then binding propagation will bind  $y$  to 3.

**Boolean Constraint Propagation:** BCP is simply arc consistency (cf. Mackworth 1992) relative to the Boolean constraint primitives `andv`, `orv`, and `notv`. BCP, like arc consistency, is semantically incomplete. For example, after evaluating (`assert! (orv z w)`) and (`assert! (orv (notv z) w)`), any variable interpretation satisfying the installed constraints must assign  $w$  the value `t`. However, BCP will not make this inference. Semantic incompleteness is necessary in order to ensure that the constraint propagation process terminates quickly. Later in the paper we discuss how we interleave backtracking search with constraint propagation to mitigate the incompleteness of local propagation.

**Generalized Forward Checking:** GFC applies

to variables for which a finite set of possible values has been established. Such a set is called an *enumerated domain*. Variables with enumerated domains are called *discrete*. For example, after evaluating (`assert! (memberv x '(a b c d))`) the variable  $x$  is discrete because its value is known to be either `a`, `b`, `c`, or `d`. Boolean variables are a special case of discrete variables where the enumerated domain contains only `t` and `nil`. Similarly, bounded integer variables are considered to be discrete. For each discrete variable SCREAMER maintains a list representing its enumerated domain. The enumerated domain for a given variable can be updated by the GFC inference process. The GFC inference process operates on constraints of the form  $z = (\text{funcall } f x_1 \dots x_n)$ . These constraints are generated by the constraint primitive `funcallv`. Unlike most constraint primitives, the primitive `funcallv` will signal an error—rather than fail—if its first argument is bound to anything but a deterministic procedure. Now consider the constraint  $z = (\text{funcall } f x_1 \dots x_n)$ . GFC will only operate on this constraint when  $f$  is bound and all but one of the remaining variables in the constraint have been bound. If the unbound variable is the output variable  $z$ , then GFC simply derives a binding for  $z$  by applying  $f$ . If the unbound variable is one of the arguments  $x_i$  then GFC tests each element  $v$  of the enumerated domain of the discrete variable  $x_i$  for consistency relative to this constraint. Elements of the enumerated domain of  $x_i$  that are inconsistent with the constraint are removed. For example, suppose that we have evaluated (`assert! (memberv x '(1 5 9))`), (`assert! (memberv y '(3 7 12))`) and (`assert! (funcallv #'< x y)`). In this case the output variable of the `funcallv` constraint is bound to `t`. Now suppose that some constraint propagation inference process infers that  $y$  is 3. In this case GFC will run on the `funcallv` constraint and remove 5 and 9 from the enumerated domain of  $x$ . Whenever the enumerated domain of a discrete variable is reduced to a single value, GFC binds the variable to that value. An example of GFC running on the  $N$ -Queens problem is given later in the paper.

**Bounds Propagation:** Bounds propagation applies to numeric variables. For each numeric value the system maintains an upper and lower bound on the possible values of that variable. These bounds propagate through constraints generated by numeric constraint primitives such as `+v`, `*v` and `<v`. For example, after evaluating (`assert! (=v z (+v x y))`), if  $z$  is known to be no larger than 5.7, and  $x$  is known to be no smaller than 2.2, then bounds propagation will infer that  $y$  is no larger than 3.5. Bounds propagation can also derive values for the Boolean output variables of numeric constraint predicates such as `<v` and `=v`. For example, if we have the constraint  $z = (< x y)$  and the system has determined that  $x$  is at least 2.0 but  $y$  is no larger than 1.0, then the system will infer that  $z$

is `nil`.

Bounds propagation will not infer a new bound unless the new bound reduces the known interval of the variable involved by at least a certain minimum percentage. This ensures that the bounds propagation process terminates fairly quickly. For example, SCREAMER avoids the very large number of bounds updates that would result from the constraints (`assert! (>v x 0)`), (`assert! (<v x 1000)`) and (`assert! (<v x (-v x 0.001))`).

**Unification:** Unification operates on constraints of the form  $w = (\text{equal } u \ v)$  which are generated by the constraint primitive `equalv`. At any given time there is a system of equations defined by the set of `equalv` constraints whose output variable has been bound to `t`. SCREAMER incrementally maintains a most general unifier  $\sigma$  for this system of equations. For example, evaluating (`assert! (equalv (cons x x) (cons y w))`) will result in a unifier  $\sigma$  that equates  $x$ ,  $y$ , and  $w$ , i.e., a unifier  $\sigma$  such that  $\sigma[x]$ ,  $\sigma[y]$ , and  $\sigma[w]$  are all equal. SCREAMER also implements disunification as in PROLOG-II (Colmerauer 1984). Thus, after evaluating (`assert! (notv (equalv x y))`), any attempt to bind  $x$  or  $y$  to be equal will fail.

The different forms of constraint propagation can interact with each other. For example, a given variable can be both discrete and numeric. The system removes non-numeric elements from the enumerated domains of discrete numeric variables. Furthermore, if a bound is known for a discrete numeric variable then elements violating that bound are eliminated from its enumerated domain. SCREAMER also derives bounds information from the enumerated domains of discrete numeric variables. Unification also interacts with SCREAMER bindings. For example, if  $\sigma$  is the most general unifier maintained by SCREAMER, and  $x$  and  $y$  are two variables such that  $\sigma[x]$  equals  $\sigma[y]$ , then any binding for  $x$  becomes a binding for  $y$  and vice versa. If  $\sigma[x]$  equals  $\sigma[y]$ , and  $x$  and  $y$  have incompatible bindings, then a failure is generated.

## Solving Systems of Constraints

By design, all of the constraint primitives described so far use only fast local propagation techniques. Such techniques are necessarily incomplete; they cannot always solve systems of constraints or determine that they are unsolvable. SCREAMER provides a number of primitives for augmenting local propagation with backtracking search to provide a general mechanism for solving systems of constraints. One such primitive, `linear-force`, can be applied to a variable to cause it to nondeterministically take on one of the values in its domain. `Linear-force` can be applied only to discrete variables or integer variables. Constraining a variable to take on a value using `linear-force` may cause local propagation. Thus a single call to `linear-force` may cause a number of variables to be

bound, or alternatively may fail if the variable cannot consistently take on any value. A second primitive, `divide-and-conquer-force`, can be applied to a variable to nondeterministically reduce the set of possible values it may take on. `Divide-and-conquer-force` can be applied only to discrete variables or real variables with finite upper and lower bounds. When applied to discrete variables, the enumerated domain is split into two subsets and the variable nondeterministically constrained to take on values from either the first or second subset. When applied to bounded real variables, the interval is split in half and the variable nondeterministically constrained to take on values in either of the two subintervals.

The above two functions operate on single variables. More generally, one must find the values of several variables which satisfy the given constraints. SCREAMER provides two primitives to accomplish this. Both are higher-order functions which take a single variable force function as an argument (e.g. `linear-force` or `divide-and-conquer-force`) and produce a function capable of forcing a list of variables using that force function. Each incorporates a different strategy for choosing which variable to force next. The first, `static-ordering`, simply forces the variables in the order given. The single variable force function is repeatedly applied to each variable, until that variable takes on a ground value, before proceeding with the next variable. All variables are bound upon termination. The second, `reorder`, selects the variable with the smallest domain, applies the single variable force function to this variable, and repeats this process until all variables are bound. Since the choice of single variable force function is orthogonal to the choice of variable ordering strategy, SCREAMER thus provides four distinct constraint solving strategies. More can easily be added.

## Examples

We will illustrate the power of the SCREAMER constraint language with two small examples. The first, shown in figure 1, solves the  $N$ -Queens problem. The function `n-queensv` creates a variable for each row and constrains each row variable to take on an integer between 1 and  $n$  indicating the column occupied by a queen in that row. The function (`a-member-ofv s`) is simply syntactic sugar for the following.

```
(let ((v (make-variable)))
  (assert! (member v s))
  v)
```

The SCREAMER primitive (`solution x f`) gathers all of the variables nested inside the structure  $x$ , applies the multiple variable forcing function  $f$  to this list of variables, and returns a copy of  $x$  where the variables have been replaced by their bound values.

In the above example, SCREAMER applies GFC as the technique for solving the underlying constraint sat-

```

(defun attacks? (qi qj distance) (or (= qi qj) (= (abs (- qi qj)) distance)))
(defun n-queensv (n)
  (solution (let ((q (make-array n)))
    (dotimes (i n) (setf (aref q i) (an-integer-betweenv 1 n)))
    (dotimes (i n)
      (dotimes (j n)
        (if (> j i) (assert! (notv (funcallv #'attacks? (aref q i) (aref q j) (- j i))))))
      (coerce q 'list))
    (reorder #'domain-size #'(lambda (x) (declare (ignore x)) nil) #'< #'linear-force)))
(defun nonlinear ()
  (let ((x (a-real-betweenv -1e40 1e40))
        (y (a-real-betweenv -1e40 1e40))
        (z (a-real-betweenv -1e40 1e40)))
    (assert! (andv (=v (+v (*v 4 x x y) (*v 7 y z z) (*v 6 x x z z)) 1356.14)
                  (=v (+v (*v 3 x y) (*v 2 y y) (*v 5 x y z)) -141.375)
                  (=v (*v (+v x y) (+v y z)) -7.7625)))
    (solution (list x y z)
              (reorder #'range-size #'(lambda (x) (< x 1e-6)) #'> #'divide-and-conquer-force))))

```

Figure 1: Two constraint-based SCREAMER programs, one for solving the  $N$ -Queens problem and one for solving a system of nonlinear equations using numeric bounds propagation.

isfaction problem. SCREAMER chooses this technique since all of the variables involved are discrete.

The second example, shown in figure 1, illustrates how bounds propagation can be used to solve systems of nonlinear equations expressed as constraints between numeric variables. The function `nonlinear` finds a solution to the following system of nonlinear equations.

$$\begin{aligned}
 4x^2y + 7yz^2 + 6x^2z^2 &= 1356.14 \\
 3xy + 2y^2 + 5xyz &= -141.375 \\
 (x + y)(y + z) &= -7.7625
 \end{aligned}$$

The expression `(a-real-betweenv -1e40 1e40)` creates a variable constrained to be a real number between the given upper and lower bounds. After the constraints have been asserted between the variables, divide and conquer search—interleaved with bounds propagation—is used to find a solution to the equations. One such solution is  $x \approx -7.311$ ,  $y \approx 6.113$ ,  $z \approx 0.367$ . Note that unlike the simplex method used in  $\text{CLP}(\mathbb{R})$ —which is limited to solving linear systems of equations—the combination of divide and conquer search interleaved with bounds propagation allows SCREAMER to solve complex nonlinear systems of equations. These techniques also enable SCREAMER to solve numeric constraint systems with both inequalities and equational constraints. Furthermore, since all of the constraint satisfaction techniques are integrated, SCREAMER can solve disjunctive systems of equations as well as systems which mix together numeric, Boolean, and other forms of constraints.

We wish to point out the intentional similarity in the names of the SCREAMER primitives `a-member-of` and

`a-member-ofv`.<sup>2</sup> Both describe a choice between a set of possible alternatives. The former enumerates that set nondeterministically by backtracking. The latter instead, creates a variable whose value is constrained to be a member of the given set. The former lends itself to a generate-and-test style of programming.

```

(let ((x1 (a-member-of s1))
      :
      (xn (a-member-of sn)))
  (unless Φ[x1...xn] (fail))
  (list x1...xn))

```

The latter lends itself to constraint-based programming.

```

(let ((x1 (a-member-ofv s1)
      :
      (xn (a-member-ofv sn)))
  (assert! (funcallv
            #'(lambda (x1...xn) Φ[x1...xn]
              x1...xn))
    (solution (list x1...xn)
              (static-ordering #'linear-force)))

```

Though these two program fragments are structurally very similar, and specify the same problem, they entail drastically different search strategies. The latter constitutes a *lifted* variant of the former. A future paper will we discuss the possibilities of performing such lifting transformations automatically. Such

<sup>2</sup>We adopt the (unenforced) convention that the names of all nondeterministic generator functions begin with the prefix `a-` or `an-` and that functions beginning with the prefix `a-` or `an-`, and also ending with `v`, denote lifted generators, functions which deterministically return a variable ranging over the stated domain instead of nondeterministically returning a value in that domain.

lifting is not limited to the `a-member-of` primitive. SCREAMER includes the following syntactic sugar for `(an-integer-between v l h)`.

```
(let ((v (make-variable)))
  (assert!
   (andv (integerpv v) (<=v v h) (>=v v l)))
  v)
```

The function `an-integer-betweenv` is a lifted analog to the SCREAMER primitive `an-integer-between`. All SCREAMER generators have lifted analogs.

## Related Work

Most of the individual techniques described in this paper are not new. What is novel is their particular combination. Programming languages which allow stating numeric constraints date back to SKETCHPAD (Sutherland 1963). Local propagation for solving systems of constraints was used by Borning (1979) in THINGLAB. Steele (1980) constructs constraint primitives very similar to ours and implements local propagation by procedural attachment. These techniques were expanded on by the MAGRITTE system (Gosling 1983). The above systems differ from SCREAMER in two ways. First, they handled only numeric constraints, lacking the GFC capacity of SCREAMER embodied in `memberv` and `funcallyv`, as well as unification and disunification embodied in `equalv`. More importantly, the constraint solving techniques incorporated in all of these systems were incomplete, particularly those based on local propagation. None of these systems could resort to interleaving backtracking search with local propagation—as SCREAMER can—to provide a slow but complete fallback to faster but incomplete local propagation techniques when applied alone.

More recently, numerous systems such as CLP( $\mathcal{R}$ ) and CHIP have been constructed in the logic programming framework which add some form of constraint satisfaction—sometimes based on local propagation—to the backtracking search mechanism already present in logic programming languages. SCREAMER differs from such systems in a number of ways, some minor and some major. First, SCREAMER uses only fast local propagation techniques as part of its constraint mechanism. The numeric constraint mechanism of CLP( $\mathcal{R}$ ) instead uses more costly techniques based on the simplex method for linear programming. These techniques are incomplete for nonlinear constraints. CLP( $\mathcal{R}$ ) and CHIP do not provide mechanisms for dealing with this incompleteness. SCREAMER, on the other hand, can solve nonlinear constraints using `divide-and-conquer-force` combined with local propagation. The second difference lies in using COMMON LISP instead of PROLOG as a substrate for constructing constraint-based programming languages. Given the substrate of nondeterministic COMMON LISP—especially its capacity for local side effects—the SCREAMER constraint package can be written totally in COMMON LISP. This gives

SCREAMER three advantages over CLP( $\mathcal{R}$ ) and CHIP. First, SCREAMER is portable to any COMMON LISP implementation. Second, SCREAMER can be easily modified and extended, to experiment with alternative constraint types and constraint satisfaction methods. Finally, SCREAMER can coexist and inter-operate with other current or future extensions to COMMON LISP such as CLOS and CLIM.

The current version of SCREAMER, including the full constraint package, is available by anonymous FTP from the file `/com/ftp/pub/screamer.tar.Z` on the host `ftp.ai.mit.edu`. We encourage you to obtain a copy of SCREAMER and give us feedback on your experiences using it.

## References

- Alan Hamilton Borning. *THINGLAB—A Constraint-Oriented Simulation Laboratory*. PhD thesis, Stanford University, July 1979. Also available as Stanford Computer Science Department report STAN-CS-79-746 and as XEROX Palo Alto Research Center report SSL-79-3.
- A. Colmerauer. Equations and inequations on finite and infinite trees. In *2d International Conference on Fifth Generation Computer Systems*, pages 85–99, 1984.
- James Gosling. *Algebraic Constraints*. PhD thesis, Carnegie-Mellon University, 1983.
- E. R. Hansen. On the solution of linear algebraic equations using interval arithmetic. *Mathematical Computation*, 22:153–165, 1968.
- Joxan Jaffar and Jean-Louis Lassez. Constraint logic programming. In *Proceedings of the 14<sup>th</sup> ACM Symposium on the Principles of Programming Languages*, pages 111–119, 1987.
- Alan K. Mackworth. Constraint satisfaction. In Stuart C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 285–293. John Wiley & Sons, Inc., New York, 1992.
- Jeffrey Mark Siskind and David Allen McAllester. SCREAMER: a portable efficient implementation of nondeterministic COMMON LISP. Technical Report IRCS-93-03, University of Pennsylvania Institute for Research in Cognitive Science, 1993.
- Ivan E. Southerland. *SKETCHPAD: A Man-Machine Graphical Communication System*. PhD thesis, Massachusetts Institute of Technology, January 1963.
- Guy Lewis Steele Jr. *The Definition and Implementation of a Computer Programming Language Based on Constraints*. PhD thesis, Massachusetts Institute of Technology, August 1980. Also available as M. I. T. VLSI Memo 80-32 and as M. I. T. Artificial Intelligence Laboratory Technical Report 595.
- Pascal Van Hentenryck. *Constraint Satisfaction in Logic Programming*. M. I. T. Press, Cambridge, MA, 1989.