Originally appeared at Supercomputing '91
This expanded version appeared in Comm. of the ACM, August 1992

The Omega Test: a fast and practical integer programming algorithm for dependence analysis

William Pugh

Dept. of Computer Science and Institute for Advanced Computer Studies
Univ. of Maryland, College Park, MD 20742
pugh@cs.umd.edu, (301)-405-2705

July 12, 1993

Abstract

The Omega test is an integer programming algorithm that can determine whether a dependence exists between two array references, and if so, under what conditions. Conventional wisdom holds that integer programming techniques are far too expensive to be used for dependence analysis, except as a method of last resort for situations that cannot be decided by simpler methods. We present evidence that suggests this wisdom is wrong, and that the Omega test is competitive with approximate algorithms used in practice and suitable for use in production compilers. Experiments suggest that, for almost all programs, the average time required by the Omega test to determine the direction vectors for an array pair is less than 500 μsecs on a 12 MIPS workstation.

The Omega test is based on an extension of Fourier-Motzkin variable elimination (a linear programming method) to integer programming, and has worst-case exponential time complexity. However, we show that for many situations in which other (polynomial) methods are accurate, the Omega test has low order polynomial time complexity.

The Omega test can be used to project integer programming problems onto a subset of the variables, rather than just deciding them. This has many applications, including accurately and efficiently computing dependence direction and distance vectors.

1 Introduction

A fundamental analysis step in an advanced optimizing compiler (as well as many other software tools) is data dependence analysis for arrays: deciding if two references to an array can refer to the same element and if so, under what conditions. This information is used to determine allowable program transformations and optimizations. For example, we can determine that in the following program, no location of the array is both read and written. Once we also verify that no location is written more than once, we know that the writes can be done in any order.

```
for i = 1 to 100 do
  for j = i to 100 do
    A[i, j+1] = A[100,j]
```

There has been extensive study of methods for deciding array data dependences [All83, BC86, AK87, Ban88, Wol89, LYZ89, LY90, GKT91, MHL91]. Much of this work has focused on approximate methods that are guaranteed to be fast but only compute exact results in certain (commonly occurring) special cases. In other situations, approximate methods are conservative: they accurately report all actual dependences, but may report spurious dependences.

Data dependency problems are equivalent to deciding whether there exists an integer solution to a set of linear equalities and inequalities, a form of integer programming. The above problem would be formulated as an integer programming problem shown below. In this example, i_w and j_w refer to the values of the loop variables at the time the write is performed and i_r and j_r refer to the values of the loop variables at the time the read is performed.

$$\begin{gathered} 1 \le i_w \le j_w \le 100 \\ 1 \le i_r \le j_r \le 100 \\ i_w = 100 \\ j_w + 1 = j_r \end{gathered}$$

Conventional wisdom holds that integer programming techniques are far too expensive to be used for dependence analysis, except as a method of last resort for situations that cannot be decided by simpler, special-case methods. We present evidence that suggests this wisdom is wrong. We describe the Omega test, which determines whether there is an integer solution to an arbitrary set of linear equalities and inequalities. We describe experiments that suggest that, for almost all programs, the average time required by the Omega test to determine the direction vectors for an array pair is less than 500 μsecs on a 12 MIPS workstation. We also found that the time required by the Omega test to analyze a problem is rarely more than twice the time required to scan the array subscripts and loop bounds. This would indicate that the Omega test is suitable for use in production compilers.

Conceptually, the Omega test combines new methods for eliminating equality constraints with an extension of Fourier-Motzkin variable elimination to integer programming. At a more detailed level, the Omega test also incorporates several implementation details (described in Section 2.4) that produce substantial speed improvements in practice.

Integer programming is a NP-Complete problem, and the Omega test has exponential worst-case time complexity. We show in Section 7 that in many situations in which other (polynomial) methods are accurate, the Omega test has low-order polynomial worst-case time complexity.

Dependence analysis is often structured as a decision problem: tests simply answer yes or no. Compilers and other program restructuring tools need to know the data dependence direction vector [Wol82] and data dependence distance vector [KMC72, Mur71] that describes the relation between the iterations in which the conflicting reads/writes occur. The data dependence distance vector describes the differences between the values of the common loop variables between the first and second access to the same array element. For example, in the following code fragment, the dependence distance of the flow dependence is (1,2):

```
for i := 1 to n do
  for j := 1 to m do
    A(i, j) := A(i-1, j-2)
```

Sometimes, dependence distance is not constant. In these cases, the dependence direction vector describes the possible combinations of signs of dependence distances.

Determining dependence direction vectors may require an exponential number calls to a dependence testing algorithm that only returns yes/no. To be competitive, a dependence analysis method must be able to short-cut this enumeration process (e.g., see [BC86, GKT91]). In Section 4, we show how the Omega test can be modified to *project* integer programming problems onto a subset of the variables, rather than just *deciding* them. With this in hand, we can efficiently produce a set of constraints that precisely and concisely describes all possible dependency distance vectors. This information can be used directly in deciding the validity of program transformations, or standard direction and distance vectors can be quickly computed from it. These techniques are described in Section 5.1.

2 The Omega test

The Omega test determines whether there is an integer solution to an arbitrary set of linear equalities and inequalities, referred to as a problem. The input to the Omega test is a set of linear equalities ($\sum_{1\leq i\leq n} a_i x_i = c$) and inequalities ($\sum_{1\leq i\leq n} a_i x_i \geq c$). To simplify our presentation (and our algorithms), we define $x_0 = 1$ and use $\sum_{0\leq i\leq n} a_i x_i = 0$ and $\sum_{0\leq i\leq n} a_i x_i \geq 0$ as our standard representations, and we use V to denote the set of indices of the variables being manipulated (i.e., $V = \{i \mid 0 \leq i \leq n\}$).

2.1 Normalizing (and tightening) constraints

Throughout this paper, we assume that any constraint we are manipulating has been normalized. A normalized constraint in one in which all the coefficients are integers and the greatest common divisor of the coefficients (not including a_0) is 1.

If the initial constraints involve rational coefficients, they can be scaled to obtain integer coefficients. (the algorithms described here do not produce any non-integer coefficients).

To normalize a constraint, we compute the greatest common divisor g of the coefficients $a_1, \ldots, a_n$. We then divide all the coefficients by g. If the constraint is an equality constraint and g does not evenly divide a_0, the constraint is unsatisfiable. If the constraint is an inequality constraint, we take the floor when dividing a_0 by g (i.e., we replace a_0 with $\lfloor a_0/g \rfloor$).

Taking floors in the constant term tightens the inequalities. If a problem P has rational but not integer solutions, tightening P may produce a problem without rational solutions, thus making it easier to determine that P has no integer solutions.

2.2 Equality constraints

Given a problem involving equality and inequality constraints, we first eliminate all the equality constraints, producing a new problem of inequality constraints that has integer solutions if and only if the original problem had integer solutions. Of course, in the process we might decide that the problem has no integer solutions regardless of the inequality constraints.

The Generalized GCD test [Banerjee88] can be used to eliminate integer equality constraints. However, we found the following approach better suited toward our needs, since it is somewhat simpler and more appropriate for situations in which additional equalities may be added later.

To eliminate the equality $\sum_{i\in V} a_i x_i = 0$, we first check if there exists a $j \neq 0$ such that $|a_j| = 1$. If so, we eliminate the constraint by solving for x_j and substitute the result into all other constraints.

Otherwise, let k be the index of the variable with the coefficient that has the smallest absolute value ($k \neq 0$) and let $m = |a_k| + 1$. We define $\widehat{\text{mod}}$ as follows:

$$a \mathbin{\widehat{\text{mod}}} b = a - b\lfloor a/b + 1/2 \rfloor$$

We create a new variable σ and produce the constraint:

$$m\sigma = \sum_{i\in V} (a_i \mathbin{\widehat{\text{mod}}} m) x_i$$

Note that $a_k \mathbin{\widehat{\text{mod}}} m = -\text{sign}(a_k)$. We then solve this constraint for x_k

$$x_k = -\text{sign}(a_k) m\sigma + \sum_{i\in V-\{k\}} \text{sign}(a_k)(a_i \mathbin{\widehat{\text{mod}}} m) x_i$$

and substitute the result in all constraints. In the original constraint, this substitution produces:

$$-|a_k| m\sigma + \sum_{i\in V-\{k\}} (a_i + |a_k|(a_i \widehat{\text{mod}} m)) x_i = 0$$

substitution	resulting constraints
Original problem	$7x + 12y + 31z = 17$ $3x + 5y + 14z = 7$ $1 \leq x \leq 40$ $-50 \leq y \leq 50$
$x = -8\sigma - 4y - z - 1$	$-7\sigma - 2y + 3z = 3$ $-24\sigma - 7y + 11z = 10$ $1 \leq -8\sigma - 4y - z - 1 \leq 40$ $-50 \leq y \leq 50$
$y = \sigma + 3\tau$	$-3\sigma - 2\tau + z = 1$ $-31\sigma - 21\tau + 11z = 10$ $1 \leq -1 - 12\sigma - 12\tau - z \leq 40$ $-50 \leq \sigma + 3\tau \leq 50$
$z = 3\sigma + 2\tau + 1$	$2\sigma + \tau = -1$ $1 \leq -2 - 15\sigma - 14\tau \leq 40$ $-50 \leq \sigma + 3\tau \leq 50$
$\tau = -2\sigma - 1$	$1 \leq 12 + 13\sigma \leq 40$ $-50 \leq -3 - 5\sigma \leq 50$
after normalization	$0 \leq \sigma \leq 2$

Figure 1: Example of elimination of equality constraints

Since $|a_k| = m - 1$, this is equal to

$$-|a_k| m\sigma + \sum_{i \in V - \{k\}} ((a_i - (a_i \widehat{\mathrm{mod}} m)) + m(a_i \widehat{\mathrm{mod}} m))x_i = 0$$

Since all terms are now divisible by m, normalizing the constraint produces:

$$-|a_k|\sigma + \sum_{i \in V - \{k\}} ((\lfloor a_i/m + \frac{1}{2} \rfloor + (a_i \widehat{\mathrm{mod}} m))x_i = 0$$

In the original constraint, the absolute value of the coefficient of σ is the same as the absolute value of the original coefficient of x_k. For all other variables, the absolute value of coefficients are reduced to at most 2/3'rds of their previous value. Therefore, repeated applications of this rule will eventually force a unit coefficient to appear and allow us to eliminate the constraint.

An application of these methods is shown in Figure 1.

2.3 Inequality constraints

The following process is used once all equality constraints have been eliminated. We first check to see if any two inequality constraints directly contradict one another (e.g., the constraints $3x + 5y \geq 2$ and $3x + 5y \leq 0$). If we find a contradiction, we report that the problem has no solutions. We can deal with equality constraints more efficiently than inequality constraints. Therefore, if we find a pair of tight inequalities (such as $6 \leq 3x + 2y$ and $3x + 2y \leq 6$), we replace them with the appropriate equality constraint and revert to our methods for dealing with equality constraints. While checking for contradictory pairs of constraints, we also eliminate constraints that are made redundant by a single other constraint (e.g., $x + 2y \geq 0$ is made redundant by $x + 2y \geq 5$).

If the problem involves at most one variable and has passed the above tests, we report that it has integer solutions. Otherwise, we reduce the problem to one or more integer programming problems in fewer dimensions and repeat the above process, eventually getting to problems in one dimension.

2.3.1 Detecting real solutions using Fourier-Motzkin variable elimination

Fourier-Motzkin variable elimination [DE73] eliminates a variable from a linear programming problem. Intuitively, Fourier-Motzkin variable elimination finds the $n-1$ dimensional shadow cast by an n dimensional object.

Consider the dodecahedron in Figure 2a. We want to calculate the shadow of the dodecahedron when it is projected along the z dimension onto the xy plane (as shown). This dodecahedron and its shadow can each be specified by a set of 12 constraints (Figure 2b).

Consider two constraints on z: a lower bound $\beta \leq bz$ and an upper bound $az \leq \alpha$ (where a and b are positive integers). We can combine these constraints to get $a\beta \leq abz \leq b\alpha$. The constraint $a\beta \leq b\alpha$ is the shadow of the intersection of these two constraints (shown visually in Figure 2c). By combining the shadow of the intersection of each pair of upper and lower bounds on z (Figures 2d and 2e), we obtain a set of constraints that defines the shadow of the original object.

Since the shadow obtained this way describes real solutions, we refer to it as the **real shadow** of a set of constraints. If there are no integer points in the real shadow of a set of constraints, we know that there are no integer solutions to the set of constraints.

Note that the set of constraints we obtained includes many redundant constraints. Performing Fourier-Motzkin variable elimination can square the number of constraints and produce many redundant constraints. However, few loop nests have dodecahedrons for iteration spaces, and in practice the number of constraints does not grow substantially. Attempting to determine which constraints are redundant so as to remove them from consideration is usually not cost-effective.

2.3.2 Detecting integer solutions using Fourier-Motzkin variable elimination

There may be integer grid points in the shadow of an object, even if the object itself contains no integer points (consider the shadow of a very thin object). Ideally, we would like to calculate the **integer shadow** of an object: a shadow such that for every integer point in the shadow, there is at least one corresponding integer point in the object above it, and vis-versa. Unfortunately, we can't always do this. The integer shadow of a convex region may not even be a convex region. However, we have developed new methods that work well in practice. Our approach can be (informally) visualized as finding the shadow of a translucent object: thicker parts of the object cast a darker shadow. If we define **dark shadows** appropriately, we can guarantee that for every integer point in the dark shadow, there is an integer point in the object above it.

As an example, we reconsider our previous example of the dodecahedron, although we squish the dodecahedron to illustrate the difficulty of finding integer shadows. The squished dodecahendron is shown in Figure 3a, and the integer points in the dodecahendron and its shadow are shown in Figures 3b and 3c. There are integer points in the real shadow that have no integer point in the object above them. However, for every integer point in the dark shadow, there is an integer point in the object above it.

The shadow is clearly dark below any part of the object that is at least one unit thick. However, since the coefficients of the constraints are integers, we can determine a looser definition of dark that will still guarantee that any integer point in the dark shadow has an integer point above it.

To determine the dark shadow, consider the case in which there is an integer solution to $a\beta \leq b\alpha$, but there is no integer solution to $a\beta \leq abz \leq b\alpha$ (i.e., there is no multiple of ab between $a\beta$ and $b\alpha$). Note that a and b are positive integers. In this case, let $i = \lfloor \beta/b \rfloor$. Then

$$ab\ i < a\beta \leq b\alpha < ab\ (i+1)$$

Since $ab(i+1) - b\alpha \geq b$, $a\beta - abi \geq a$ and $ab(i+1) - abi = ab$, $b\alpha - a\beta \leq ab - a - b$. If $b\alpha - a\beta \geq ab - a - b + 1 = (a-1)(b-1)$, we know that there must be an *integer* solution to z. Therefore, the dark shadow of $\alpha \geq az$ and $bz \geq \beta$ is:

$$b\alpha - a\beta \geq (a-1)(b-1)$$

Note that if $a = 1$ or $b = 1$, the dark shadow and the real shadow are identical. If the dark and real shadow resulting from each combination of an upper and lower bound are identical, the projection is called an *exact* projection. This will happen, for example, if all of the coefficients of z in lower bounds on z are 1, or if all of the coefficients of z in upper bounds on z are 1. For the problems that arise in dependence analysis, we can almost always find an exact projection.

We now have a method for checking for the existence of integer solutions to a set of constraints:

1. We first decide which variable to eliminate. We choose this variable so as to perform an exact projection if possible, and to minimize the number of constraints resulting from the combination of upper and lower bounds. If we are forced to perform non-exact reductions, we choose a variable with coefficients as close to zero as possible.
2. Calculate the real and dark shadows of the set of constraints along that dimension.
3. If the real and dark shadows are identical, there are integer solutions to the original set of constraints iff there are integer solutions to the shadow.
4. Otherwise:
   (a) If there are no integer solutions to the real shadow, we know there are no integer solutions to the original set of constraints.
   (b) If there are integer solutions to the dark shadow, we know there are integer solutions to the original constraints.
   (c) Otherwise, we know if an integer solution exists, it must be closely nestled between an upper bound and a lower bound. Therefore we consider a set of planes that are parallel to a lower bound and close to a lower bound. Any integer solution closely nestled between an upper bound and a lower bound must lie on one of these planes.

   Computationally, we analyze the problem as follows:

   We know that if there exists an integer solution to the original set of constraints, there must exist a pair of constraints $\alpha \geq az$ and $bz \geq \beta$ on z such that

$$ab - a - b \geq b\alpha - a\beta \geq 0 \wedge b\alpha \geq abz \geq a\beta$$

$$ab - a - b + a\beta \geq abz \geq a\beta$$

   We check this by determining the largest coefficient a of z in any upper bound on z, and, for each lower bound $bz \geq \beta$ on z, testing if there are integer solutions to the original problem combined with $bz = \beta + i$ for each i such that $(ab - a - b)/a \geq i \geq 0$.

   While these steps are expensive and complicated, they rarely, if ever, need to be used in practice.

2.3.3 An Omega test nightmare

To demonstrate (and show the limitations of) the techniques used, we illustrate the steps performed by the Omega test on an example designed to force the Omega test to work very hard for a small problem. Consider the inequalities P:

$$\begin{array}{rcl} 27 \leq & 11x + 13y & \leq 45 \\ -10 \leq & 7x - 9y & \leq 4 \end{array}$$

There are no exact projections we can perform, and we would decide to eliminate x since the coefficients of x are (slightly) smaller. Figure 4a shows the constraints in the original problem, and the unnormalized constraints in the real and dark shadows. Since the real shadow has integer solutions but the dark shadow does not, we check if there are any integer solutions close to a lower bound. We do this by checking if the intersection of the original set of constraints and any one of the following constraints contains an integer point (this is shown graphically in Figure 4b). Since there are no such solutions, we know that no integer solutions exist.

$$\begin{array}{rl} 7x = 9y - 10 + j & 0 \leq j \leq \lfloor \frac{77-11-7}{11} \rfloor = 5 \\ 11x = 27 - 13y + j & 0 \leq j \leq \lfloor \frac{121-11-11}{11} \rfloor = 9 \end{array}$$

The steps performed in this example appear complicated and expensive. However, this example was *designed* to be expensive to resolve. We do not expect situations this difficult to arise frequently in practice. Also, although many steps are performed in this process, our implementation of the Omega test takes only 4.5 milliseconds on a 12 MIPS workstation to perform them all.

Worse nightmares are possible: on problems with only 2 variables and 3 constraints, the Omega test can take time proportional to the absolute value of the coefficients. While this is a frightening possibility, we do not expect these situations to arise frequently in practice.

A decision on better methods for dealing with Omega test nightmares will have to wait until more experience is gained about the type of nightmares that occur in practice.

2.4 Implementation Details

In implementing the Omega test we used several algorithmic ideas and tricks that substantially improved our running time. We report some of those ideas here.

Equalities and inequalities are represented as vectors of coefficients. The Omega test is crafted so that the algorithms only need to deal with integers; no rational number representation scheme needs to be used.

Once we have eliminated all the equality constraints from a problem, we check for any variables that have no lower bounds or have no upper bounds. We refer to such variables as unbounded variables. Performing Fourier-Motzkin elimination on an unbounded variable simply deletes all the constraints involving it. We delete all constraints involving unbounded variables and then check if that has produced additional unbounded variables. We repeat this process until no unbounded variables remain.

Next, we normalize all the constraints and assign hash keys and constraint keys to them. We only do this to constraints that have been modified since the last time they were normalized. The constraint key of a constraint is a unique tag based on the coefficients of the variables in the constraint; two constraints have equal constraint keys if and only if they differ only in their constant term. Constraint keys are both negative and positive, and the key of a constraint e_1 is the negation of the key of a constraint e_2 if and only if the coefficients of the variables in e_1 are the negation of the coefficients of the variables in e_2. We refer to this as opposing keys and opposing constraints. Constraint keys are assigned to constraints in constant expected time by recording in a hash table constraint keys previously assigned. We compute a hash key based on the coefficients of the constraint as an index into the hash table (hash keys are not guaranteed to be unique). Our method for computing hash keys is designed so that opposing constraints have opposing hash keys, which makes it easy to assign them opposing constraint keys. As constraints are normalized, we enter them into a table based on their constraint key. This allows us to check for redundant, contradictory or tight constraint pairs in constant time per constraint.

In the process of normalizing constraints, we check to see if any constraints involve more than one variable. After normalization, if we found no multi-variable constraints, we know the system must have solutions, and we return immediately.

Next, we examine the variables to decide which variable to eliminate. If we can perform an exact projection, we perform the elimination in place (adding and deleting constraints from the current problem). Otherwise, we copy the constraints with zero coefficients for the eliminated variable into two new problem data structures (for the real shadow and for the dark) and then add the constraints produced by Fourier-Motzkin elimination. Since the constraints generated for the real and dark shadow differ only in their constant terms, we can share much of the work in adding these constraints.

3 Nonlinear subscripts

Integer programming dependence analysis methods allow us to properly handle symbolic constants [LT88, HP90] and some types of `min` and `max` functions in loop bounds [WT92] and conditional assignments [LC90].

For example, even if we had no information about the value of `n`, we would like to be able to decide that there are no flow dependences in the following program:

```
for i = 1 to n do
   a[i+n] = a[i]
```

As previous authors have suggested, we can handle loop-invariant symbolic constants by adding them as additional variables to the integer programming problem. For example, the above problem would generate the following integer programming program (involving the variables i, i' and n):

$$1 \leq i, i' \leq n$$
$$i + n = i'$$

We also can accommodate integer division and integer remainder operations, something that does not appear to have been previously recognized. Assume an expression `e` appears in a program that can be expressed as $\mathtt{e} = \alpha \ \mathtt{div}\ m$ where m is a positive integer. To handle this, we define a new variable σ, add the inequality constraints $0 \leq \alpha - m\sigma \leq m - 1$ and use σ as the value of `e`. Similarly, if $\mathtt{e} = \alpha \ \mathtt{mod}\ m$ we would add the same inequality constraint but use $\alpha - m\sigma$ as the value of `e`.

4 Projection of Integer Programming Problems

As described in Section 2, the Omega test simply *decides* if there is a solution to an integer programming problem. In this section, we describe how to adapt the Omega test to allow it to be used for *symbolic projection*. When used this way, the Omega test is given as input an integer programming program P and a designation of a set of protected variables $\widehat{V} \subset V$. The Omega test projects P into one or more problems involving only variables in $\widehat{V}$ that describe all the possible values of the variables in $\widehat{V}$ such that there is an integer solution to P with those values. For example, projecting the integer programming problem $\{0 \leq a \leq 5;\ b < a \leq 5b\}$ onto a produces the problem $\{2 \leq a \leq 5\}$.

Actually, results of the projection process can be a little more complicated than just described. The results may not be in terms of the variables in $\widehat{V}$. Instead, the results are given in terms of a set $\widehat{V}'$ of not more than $|\widehat{V}|$ variables (possibly including new variables), along with methods for calculating the appropriate values for the values of $\widehat{V}$ from the values of $\widehat{V}'$. For example, if asked to project the integer programming problem $\{a = 10b + 25c;\ a \geq 13\}$ onto a, the Omega test will produce $\{\sigma \geq 3;\ a = 5\sigma\}$.

The projection process may produce multiple problems. For example, projecting the problem $\{5b \leq a \leq 6b\}$ onto a produces:

$$\{20 \leq a\}$$
$$\{0 \leq \sigma;\ a = 6\sigma\}$$
$$\{1 \leq \sigma;\ a = 6\sigma - 1\}$$
$$\{2 \leq \sigma;\ a = 6\sigma - 2\}$$
$$\{3 \leq \sigma;\ a = 6\sigma - 3\}$$

4.1 Changes to the Omega test

Three of the changes required are simple, the other is not as simple. The quick changes are:

- If the current problem P involves only protected variables, check if there are integer solutions of P and if so, report P as one projection.

- When performing an inexact Fourier-Motzkin elimination, project the dark shadow and the intersection of the original constraints with all of the equality checks near the lower bounds. In other words, we must project all of the subproblems where we would look for an integer solution, not stopping when an integer solution is first verified. This could be expensive if projecting a system involves many inexact projections. We do not believe this will occur in practice for the problems arising from dependency analysis.
- We never perform Fourier-Motzkin variable elimination on a protected variable. This could require us to perform a inexact projections in a situation where we could have performed an exact projection if we were not protecting certain variables.

The not so simple change involves equalities. Given an equality constraint $\sum_{i \in V} a_i x_i = 0$, let g be the gcd of the coefficients of the non-protected variables. (we assume (as always) that the constraint is normalized).

- If $g = 0$, the constraint involves only protected variables. We use our standard methods to eliminate the constraint. This will result in the elimination of a protected variable. All substitutions performed in this process are recorded in a substitution log. These substitutions involve only protected variables.
- If $g = 1$, we use our standard techniques (Section 2.2) to find a substitution involving only unprotected variables that simplifies or eliminates the constraint.
- If $g > 1$, we create a new protected variable σ, add the constraint:

  $$g\sigma = \sum_{i \in V} (a_i \ \widehat{\bmod}\ g) x_i$$

  Eliminating this new constraint will transform the original constraint so that the gcd of the non-protected variables is 1 (after normalization).

When we report a projection, any substitutions involving protected variables are translated back into equality constraints.

4.2 Projection with wildcards

As a modification of the approach described above, we could refuse to perform inexact reductions while performing projection. The advantage of this is that we only report one projected problem as our result. The disadvantage is that the projected problem has additional variables (that should be treated as wildcards)

In the applications we have found for projection, we have found projection with wildcards to be more useful than producing multiple results.

5 Using projection

This projection technique can be used for several purposes. We describe some that have occurred to us.

5.1 Dependence direction and distance vectors

One problem with some dependence analysis methods is that they are only "yes/no" decision methods. In compilers and other program structuring tools, we need to know the data dependence direction vector [Wol82] and data dependence distance vector [KMC72, Mur71] describing the relation between the iterations in which the conflicting reads/writes occur. One way to determine dependence direction vectors is to make 3^L calls to a decision procedure (where L is the number of loops surrounding both references). In order to be competitive, a dependence analysis method must be able to short-cut this enumeration (for example, see [BC86, GKT91]).

In our method, we take the integer programming problem for determining if any dependence exists between two references, and introduce a new variable for the dependence distance in each shared loop (along with the appropriate equality constraints to define the value of the variable). We then project the problem onto the dependence distance variables. The projected system may be a better way to describe dependence conditions than dependence directions and distances; it accurately describes more information than is typically contained in dependence direction vectors (such as when a dependence distance is always greater than 5).

Alternatively, we can use the projected set of constraints to determine efficiently the dependence direction and distance vectors. We scan the dependences, and infer as much information as possible from constraints involving a single dependence distance variable. We next unprotect any dependence distance variable that is uncoupled or who's sign is completely determined. If coupled variables were unprotected, we project the problem onto the protected variables and repeat this process. Otherwise, we choose one protected variable and generate the subproblems for two or three possible signs for the variable (negative, zero or positive), and recursively explore those.

For example, the dependence distances for the following array pair

```
for j = 0 to 20 do
  for i = max(-j,-10) to 0 do
    for k = max(-j,-10)-i to -1 do
      for l = 0 to 5 do
        a(l,i,j) = ...
              ... = a(l,k,i+j)
```

simplify to:

$$
\begin{aligned}
&0 \leq \Delta j \leq 10\\
&\Delta i + \Delta j \leq 10\\
&\Delta i + 2\Delta j \leq 10\\
&3\Delta j + 2\Delta i + \Delta k \leq 20\\
&2\Delta j + 2\Delta i + \Delta k \leq 10\\
&1 \leq \Delta j + \Delta i + k\\
&1 \leq \Delta j + \Delta i\\
&\Delta l = 0
\end{aligned}
$$

We first unprotect Δl, and then consider $\text{sign}(\Delta j) = 0$ and $\text{sign}(\Delta j) = 1$. Considering $\text{sign}(\Delta j) = 0$ gives:

$$1 \leq \Delta i \leq 10; 1 \leq \Delta k + \Delta i; \Delta k + 2\Delta i \leq 10$$

We would then unprotect Δi (since we know $\text{sign}(\Delta i) = 1$) and project the problem, obtaining $-8 \leq \Delta k \leq 8$, which gives a direction vector of $(=, <, *, =)$.

Returning to consideration of $\text{sign}(\Delta j) = 1$ produces:

$$
\begin{aligned}
&-8 \leq \Delta i \leq 8\\
&-8 \leq \Delta k \leq 8\\
&-8 \leq 2\Delta k + \Delta i\\
&-9 \leq \Delta k + \Delta i\\
&\Delta k + 2\Delta i \leq 8\\
&\Delta k - \Delta i \leq 17
\end{aligned}
$$

Recursively analyzing the possibilities for the sign of Δi produces direction vectors of $(<, >, *, =)$, $(<, =, *, =)$ and $(<, <, *, =)$. This example is the most difficult example seen in our testing, requiring 2492 μsecs to analyze.

5.2 Run-time checks and Compile-time assertions

By projecting a problem onto the variables corresponding to symbolic constants that cannot be determined at compile-time, we can produce a predicate that will allow us to determine at run-time if a particular dependence or dependence direction exists (as described by [KK91]). Alternatively, at compile time we could ask the user if the predicate is true.

5.3 Summarizing Array References

In interprocedural analysis, we need to characterize the portions of an array that may be affected by a procedure call [Tri85, BK89, HK90, IJT91]. We can use the Omega test to obtain an accurate summary of the locations of an array that might be affected by a single assignment statement. We do this by setting up an integer programming problem involving variables for each array index and all loop variables and symbolic constants, and adding appropriate constraints for the loop bounds, subscript expressions, and so on. Projecting this problem onto the variables for the array indices and the symbolic constants gives an accurate summary of the locations of the array affected by the assignment statement. The summary is not limited to convex polyhedron. The projected problem will have solutions only for those locations that can actually be changed. Details such as strides are accurately represented.

The Omega test can easily be used to determine when two regions intersect. With more work, the Omega test can be used to check if one region is a subset of another. It is unclear how to use the Omega test to merge affected regions; however, the Omega test could be used to convert exact affected regions into approximate affect regions (such as described by [BK89, HK90]) and then those regions could be merged.

5.4 Determining Loop Bounds

The Omega test can be used to determine appropriate loop bounds when interchanging non-rectangular loops. This use of integer programming and projection to perform this is described by [AI91].

6 Performance

We have implemented the Omega test in Wolfe's `tiny` tool [Wol91]. We handle `min` and `max` expressions in loop bounds and symbolic constants, and compute exact sets of direction vectors (as opposed to the compressed direction vectors normally generated by `tiny`). We applied this tool to the programs 1, 3, 4, 5 and 7 of the NASA NAS benchmark suite and to all the `tiny` source files distributed with `tiny`, (which include Cholesky decomposition, LU decomposition, several versions of wavefront algorithms, and several more contrived examples), as well as several of our own test programs. Programs 2 and 6 of the NAS benchmark make extensive use of index arrays. Since we do not provide special treatment for index arrays, we decided that it would be misleading to include them. The analysis of array pairs that have different constant subscripts (e.g., `a(4)` and `a(5)`) are *not* included in the figures reported here; those cases are detected while scanning the subscripts (thus both avoiding the analysis time and the time required to scan the loop bounds). Standard optimizations such as induction variable recognition and forward substitution were performed by hand. We did not compute input dependences (an input dependence is a dependence between two reads of the same location of an array) or dependences between array pairs that did not share at least one common loop.

We timed the Omega test on a Decstation 3100, a 12 MIPS workstation based on a MIPS R2000 CPU. Shown below are our results on the time per array pair required to analyze programs in the NASA NAS benchmark:

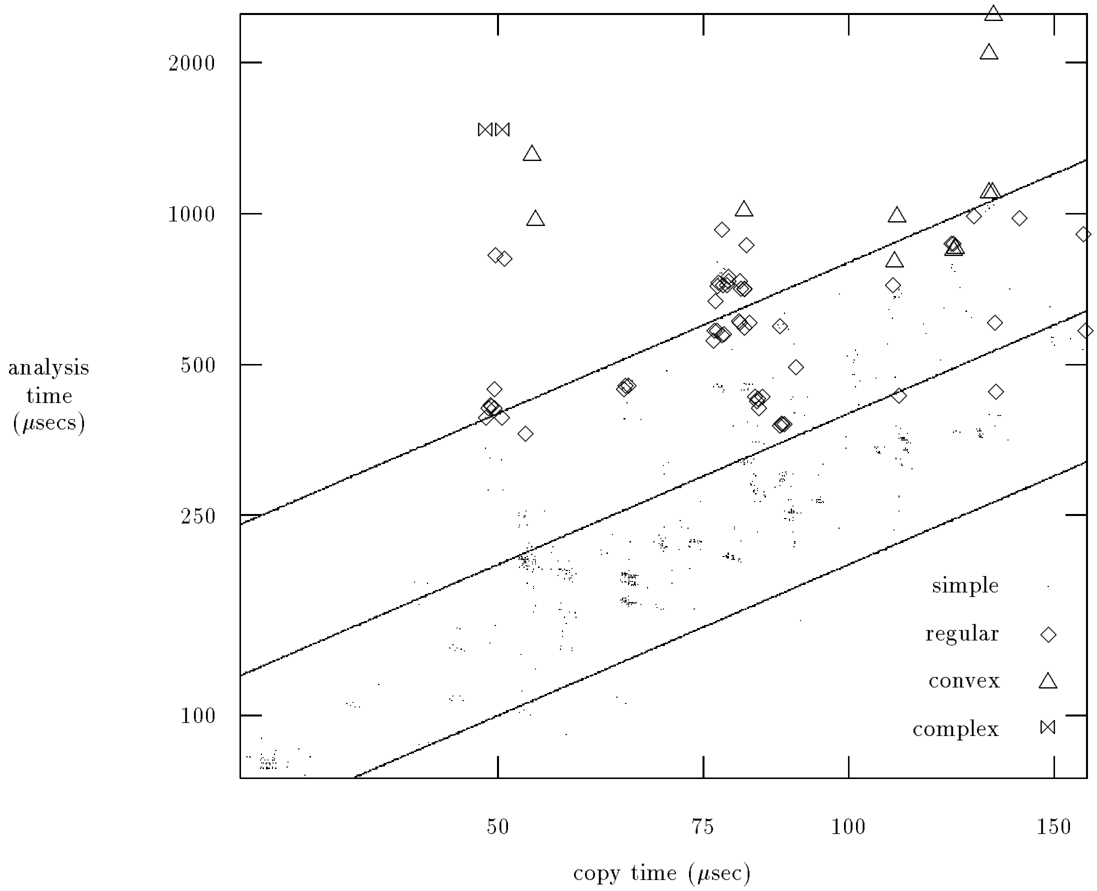

Figure 2: Omega Test Performance

Program	average time	95%-tile time
#1: MXM	275 μsecs	316 μsecs
#3: CHOLSKY	504 μsecs	1024 μsecs
#4:BTRIX	250 μsecs	367 μsecs
#5:GMTRY	191 μsecs	534 μsecs
#7:VPENTA	129 μsecs	204 μsecs

The third program of the NAS benchmark (CHOLSKY) is substantially more complicated that almost all real-world FORTRAN code, involving loops nested four deep, triplely subscripted arrays and groups of 3 coupled loop indices. We feel confident that it represents a good "worst-case example" for analyzing dusty deck FORTRAN code (excluding treatment of index arrays).

Our results on individual array pairs from all programs tested are shown in Figure 2. Each point is the timing result for a single array pair. To present the results in a somewhat machine independent fashion, the results are plotted on a log/log graph of analysis time vs. copying time (the time required just to copy the problem). All times were randomly perturbed by $\pm 1/2\mu$sec to spread out overlapping points. The diagonal lines are drawn at analysis time = 8 $\times$ copying time, 4 $\times$ copying time and 2 $\times$ copying time.

The analysis time is the total time required to analyze the array pair, calculate the appropriate direction vectors and add the dependences to dependence graph, excluding the time required to scan the array subscripts and loop bounds and build the constraints that describe the dependence between the array pairs.

Across a range of test programs, we found the following break-down for how time was spent by the Omega test: about 1/2 the time was spent dealing with inequality constraints, about 1/4 of the time was spent on dealing with equality constraints, and 1/4 of the time was spent examining projected constraints to construct direction vectors. None of our test cases required inexact Fourier-Motzkin variable elimination.

To analyze our results, the set of constraints describing the dependence distances for each array pair were analyzed to remove any redundant constraints (this is not cost-effective normally). Based on the simplified constraints, each array pair was classified as follows:

simple Any case that does not involve coupled dependence distances.

regular A case where dependence distances are coupled, but all inequality constraints have unit coefficients (for example, $\{\Delta i \geq 0; \Delta i + \Delta j > 0\}$).

convex A case where the inequality constraints define a convex region but at least one constraint has a non-unit coefficient (for example, $\{0 \leq \Delta j \leq 10; 0 \leq \Delta i + \Delta j \leq 10; \Delta i + 2\Delta j \leq 10\}$ – the last constraint makes this non-regular).

complex A case where the inequality constraints define a non-convex region. We only encountered two such cases, one shown below and another one identical except that the lower bound of the `i` loop is 2.

```
for i = 1 to 10 do
  for j = 0 to 4 do
    a(i-j) = a(j)
  endfor
endfor
```

The flow/anti dependence distances for the example above are all the distances that satisfy $\{-4 \leq \Delta j \leq 4; -7 \leq \Delta i - \Delta j, \Delta i + \Delta j \leq 10; \Delta i \leq 9\}$ except for $\{\Delta i = 9; \Delta j = 0\}$.

Maydan, Hennessy and Lam [MHL91] use memoization to obtain better performance. Memoization could be added to the Omega test. However, the cost of computing a hash key and verifying a cache hit would be about 2-4 times the copying cost for a problem, and therefore adding caching to the Omega test would not produce significant savings for typical, simple cases and may produce little or no overall speed improvement.

We found that the cost of scanning array subscripts and loop bounds to build a dependence problem was typically 2-4 times the copying cost for the problem. Thus, for many array pairs the cost of building the dependence problem was nearly as large or even larger than the time spent analyzing the resulting problem. We have not spent much effort trying to improve the performance of the code that builds dependence problems. However, it is difficult to imagine building a dependence problem in much less than twice the time required to copy the problem. This suggests that for the majority of array pairs, using a dependence analysis algorithm significantly faster than the Omega test would not lead to significant overall speed improvements.

7 Polynomial time bounds

We first describe some general time bounds on parts of the Omega test, and then describe polynomial time bounds for cases where other polynomial time algorithms are accurate. In this section, we use m to denote the number of constraints and n to denote the denote the number of variables.

The time taken by the methods in Section 2.2 to eliminate one equality constraint is $O(mn \log |C|)$ worst-case time, where C is coefficient with the largest absolute value in the constraint. This cost arises from the fact that we might have to apply the perform $\log |C|$ substitutions before we can eliminate the constraint, and performing a substitution takes $O(nm)$ time.

Eliminating unbound variables takes $O(mnp)$ worst-case time, where p is the number of passes required to eliminate all the variables that become unbound. At least one variable is eliminated in each pass except the last.

Normalizing the constraints and checking for directly contradictory or redundant constraints requires $O(mn)$ expected time (the time bound is only expected, not worst-case, because hashing is used).

Producing the subproblems resulting from Fourier-Motzkin variable elimination takes time proportional to the size of the subproblems produced.

7.1 Special cases

During normalization, the Omega test checks to see if any variables are involved in constraints with other variables. If not (and checking for contradictory constraint pairs has not produced a contradiction), we know the problem has solutions and do not need to perform any additional computation. This applies iff the "Single Variable Per Constraint" (SVPC) test [MHL91] can be applied, which was found [MHL91] to be applicable in 1/3 of the unique cases found in the Perfect Club Benchmark (a higher percentage if duplicate cases were considered separately).

The "Acyclic Test" [MHL91] can be applied in exactly those cases that the Omega test can resolve just by eliminating unbound variables and performing exact projections that do not increase the number of constraints, a process that takes $O(mn^2)$ worst-case time. They found [MHL91] that this test could be applied in over 1/4 of the unique cases encountered.

The "Loop Residue" algorithm [Sho81] can be applied in just those cases where each constraint is of the form $x_i \geq x_j + c$, $x_i \geq c$, or $c \geq x_i$. In a set of constraints with this property, Fourier-Motzkin variable elimination is exact and preserves this property. On n variables, there can be at most $n^2 + n$ constraints of this form after eliminating redundant pairs. Thus, the Omega test will take $O(n^3)$ time to resolve a set of constraints that can be solved by the Loop Residue algorithm. Maydan, Hennessy and Lam [MHL91] found that the Loop Residue algorithm could be applied in 1/4 of the unique cases encountered in their study of the Perfect Club benchmark.

Maydan, Hennessy and Lam found that 91% of the cases they encountered could be determined by constant tests and Banerjee's Generalized GCD tests. Of the remaining 9% of the cases, they found that their SVPC, Acyclic or Loop Residue tests could be applied in 86% of the unique cases.

The Delta test [GKT91] works by searching for dependence distances that can be easily determined, and then propagating that information with the intent of making it possible to easily determine other dependence distances precisely. In the cases where their algorithm can determine a dependence distance without the use of MIV tests, the Omega test also will determine it efficiently (and in polynomial time) by a combination of solving equality constraints, tightening inequality constraints and converting tight inequality constraints into equality constraints. Since the Omega test treats the dependence analysis problem as a single integer programming problem, it automatically achieves the propagation effects of the Delta test. Therefore, any dependence analysis problem that can be solved by the Delta test without resorting to exponential algorithms or approximate methods (i.e., resorting to what they refer to as MIV tests) can be solved in polynomial time by the Omega test.

In their study of the RiCEPS, Perfect, SPEC benchmarks and LINPACK and EISPACK, they found that 97% percent of the cases could be solved without requiring the use of MIV tests.

Since the Omega test can solve effectively and in (effective) polynomial time any problem that be solved by any combination of the Single Variable Per Constraint test, the Acyclic test, the Loop Residue test and the Delta test, we expect that it should be able to solve more problems exactly and efficiently than any one of them alone.

8 Related work on Exact Dependence Analysis

The Constraint-Matrix test [Wal88] makes use of the simplex algorithm modified for integer programming. The Constraint-Matrix test can fail to terminate and it is not clear how efficiently it works in practice.

Lu and Chen describe [LC90] an integer programming algorithm for dependence analysis. However, their method appears prohibitively expensive for use in a production compiler.

Triolet [Tri85] used Fourier-Motzkin techniques for representing affected array regions in interprocedural analysis. Triolet found Fourier-Motzkin techniques to be expensive (22 to 28 times longer than using simpler methods for representing affected array regions).

Several implementations of Fourier-Motzkin variable elimination have been described for use in dependence analysis. The Power test described by Wolfe and Tseng [WT92] combines the Banerjee's Generalized GCD test, constraint tightening, and Fourier-Motzkin variable elimination. They take no special action when performing an inexact projection except to flag the result as possibly being conservative. Fourier-Motzkin elimination is used by Maydan, Hennessy and Lam [MHL91] if none of the other methods they use give an exact answer. They use back substitution to determine a sample solution. If the sample solution is not integral, they suggest the use of branch and bound methods to verify or disprove the existence of integer solutions (they have not found the need to implement this thus far). Both Wolfe and Tseng [WT92] and Maydan, Hennessy and Lam [MHL91] suggest that due to the expense of Fourier-Motzkin variable elimination, simpler tests should be used instead in situations where they are known to be accurate.

Ancourt and Irigoin [AI91] describe the use of Fourier-Motzkin variable elimination so as to determine loop bounds for iterating over an iteration space described by a set of linear inequalities (using projection as described in Section 4). Their work has significant overlaps with ours.

When performing what is apparently an inexact projection, they first perform a more elaborate process to check if the projection is indeed inexact. They consider a concept similar to our dark shadow, except they use force the difference between the upper and lower bounds to be at least $(a-1)b$, as opposed to $(a-1)(b-1)$. Since our definition is safe and makes the dark shadow larger, it is the preferred choice.

They do not actually generate the dark shadow as a separate problem. Rather, they check to see if the constraints in the dark shadow are redundant with respect to the real shadow. If they are, then the dark shadow and real shadow are identical, and the elimination is exact.

If the projection is not exact, then they add *pseudo-linear* constraints to the real shadow to obtain the integer shadow. These pseudo-linear constraints appear useful and appropriate for determining loop bounds. However, they are difficult to use for determining the existence of integer solutions.

Ancourt and Irigoin [AI91] do not give any performance data for their algorithm.

A recent report [IJT91] on the PIPS project mentions that Fourier-Motzkin variable elimination is used to analyze dependences (based on the work described in [AI91]). The methods used are not fully described, but the basic framework appears similar to that described in Section 5.1. It is not clear how the pseudo-linear constraints of [AI91] are handled. They point out that in many simple cases, Fourier-Motzkin variable elimination is fast and efficient. They state that using integer programming techniques for dependence analysis incurs a very high cost (that is acceptable since PIPS is not a production system). They also state that in their implementation dependence testing does not take a noticeable amount of time compared with the whole parallelization process.

9 Source code availability

A `C` language implementation of the Omega test is freely available for anonymous `ftp` from `ftp.cs.umd.edu` in directory `pub/omega`. Files available include a stand-alone version of the Omega test and a version of Wolfe's `tiny` tool [Wol91] extended to use the Omega test.

10 Conclusions

Conservative dependence analysis methods may be efficacious for the demands of vectorizing compilers. Transforming programs so as to make efficient use of massively parallel SIMD computers is a much more demanding task. Also, programs that have undergone transformations such as loop skewing and loop interchange present analysis problems substantially more difficult than encountered in typical dusty-deck FORTRAN.

Our studies have convinced us that the Omega test is a fast and practical method for performing data dependence analysis that is not only adequate for problems encountered in vectorizing FORTRAN code, but also for the demands of more sophisticated program transformation tools.

Performing projection of integer programming problems is an exciting concept. We have discussed how it can be used to determine efficiently information about dependence direction and distance vectors, as well for several other uses. It can make it much easier to describe and build program analysis and transformation tools. For example, it can be used for determining loop bounds after loop interchange [AI91], and we have made extensive use of it in work that considers loop transformations in a uniform manner [Pug91].

11 Acknowledgements

Thanks to everyone who gave me feedback on this work, especially Michael Wolfe and the anonymous referee who provided detailed comments, as well as to my research group (Dave Wonnacott, Udayan Borkar and Wayne Kelly). This work is supported by NSF grants CCR-8908900 and CCR-9157384 and a Packard Fellowship.

12 Author's Biography

Dr. William Pugh received his Ph.D. from Cornell in 1988. He has done research in the areas of incremental computation, randomized data structures, implementation of multiple inheritance, programming languages for hard real-time systems, and compilers for supercomputers. In 1991, he received a Presidential Young Investigator Award and a Packard Fellowship. Dr. Pugh is currently an assistant professor at the University of Maryland, College Park, in the Department of Computer Science and Institute for Advanced Computer Studies.

13 CR Categories

D.3.4 Optimizing Compilers

References

[AI91] Corinne Ancourt and François Irigoin. Scanning polyhedra with do loops. In *Proc. of the Third ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 39–50, April 1991.

[AK87] J. R. Allen and K. Kennedy. Automatic translation of Fortran programs to vector form. *ACM Transactions on Programming Languages and Systems*, 9(4):491–542, October 1987.

[All83] J. R. Allen. *Dependence Analysis for Subscripted Variables and Its Application to Program Transformations.* PhD thesis, Rice University, April 1983.

[Ban88] U. Banerjee. *Dependence Analysis for Supercomputing.* Kluwer Academic Publishers, Boston, MA, 1988.

[BC86] M. Burke and R. Cytron. Interprocedural dependence analysis and parallelization. In Proceedings of the SIGPLAN '86 Symposium on Compiler Construction, Palo Alto, CA, July 1986.

[BK89] V. Balasundaram and K. Kennedy. A technique for summarizing data access and its use in parallelism enhancing transformations. In *SIGPLAN Conference on Programming Language Design and Implementation, '89*, June 1989.

[DE73] G.B. Dantzig and B.C. Eaves. Fourier-Motzkin elimination and its dual. *Journal of Combinatorial Theory (A)*, 14:288–297, 1973.

[GKT91] G. Goff, Ken Kennedy, and Chau-Wen Tseng. Practical dependence testing. In *ACM SIGPLAN'91 Conference on Programming Language Design and Implementation*, 1991.

[HK90] Paul Havlak and Ken Kennedy. Experience with interprocedural analysis of array side effects. In *Supercomputing '90*, 1990.

[HP90] M. Haghighat and C. Polychronopoulos. Symbolic dependence analysis for high performance parallelizing compilers. In *Proceedings of the Third Workshop on Languages and Compilers for Parallel Computing*, August 1990.

[IJT91] François Irigoin, Pierre Jouvelot, and Rémi Triolet. Semantical interprocedural parallelization: An overview of the pips project. In *Proc. of the 1991 International Conference on Supercomputing*, pages 244–253, June 1991.

[KK91] D. Klappholz and X. Kong. Extending the Banerjee-Wolfe test to handle execution conditions. Technical Report 9101, Dept. of EE/CS, Stevens Institute of Technology, 1991.

[KMC72] D. Kuck, Y. Muraoka, and S. Chen. On the number of operations simultaneously executable in FORTRAN-like programs and their resulting speedup. *IEEE Transactions on Computers*, 1972.

[LC90] L. Lu and M. Chen. Subdomain dependence test for massive parallelism. In *Proceedings of Supercomputing '90*, New York, NY, November 1990.

[LT88] A. Lichnewsky and F. Thomasset. Introducing symbolic problem solving techniques in the dependence testing phases of a vectorizer. In *Proceedings of the Second International Conference on Supercomputing*, St. Malo, France, July 1988.

[LY90] Z. Li and P. Yew. Some results on exact data dependence analysis. In D. Gelernter, A. Nicolau, and D. Padua, editors, *Languages and Compilers for Parallel Computing.* The MIT Press, 1990.

[LYZ89] Z. Li, P. Yew, and C. Zhu. Data dependence analysis on multi-dimensional array references. In *Proceedings of the 1989 ACM International Conference on Supercomputing*, June 1989.

[MHL91] D. E. Maydan, J. L. Hennessy, and M. S. Lam. Efficient and exact data dependence analysis. In *ACM SIGPLAN'91 Conference on Programming Language Design and Implementation*, pages 1–14, June 1991.

[Mur71] Y. Muraoka. *Parallelism Exposure and Exploitation in Programs.* PhD thesis, Dept. of Computer Science, University of Illinois at Urbana-Champaign, February 1971.

[Pug91] William Pugh. Uniform techniques for loop optimization. In *1991 International Conference on Supercomputing*, pages 341–352, Cologne, Germany, June 1991.

[Sho81] R. Shostak. Deciding linear inequalities by computing loop residues. *Journal of the ACM*, 28(4):769–779, October 1981.

[Tri85] R. Triolet. Interprocedural analysis for program restructuring with Parafrase. CSRD Rpt. 538, Dept. of Computer Science, University of Illinois at Urbana-Champaign, December 1985.

[Wal88] D. Wallace. Dependence of multi-dimensional array references. In *Proceedings of the Second International Conference on Supercomputing*, St. Malo, France, July 1988.

[Wol82] M. J. Wolfe. *Optimizing Supercompilers for Supercomputers*. PhD thesis, Dept. of Computer Science, University of Illinois at Urbana-Champaign, October 1982.

[Wol89] Michael Wolfe. *Optimizing Supercompilers for Supercomputers*. Pitman Publishing, London, 1989.

[Wol91] Michael Wolfe. The tiny loop restructuring research tool. In *Proc of 1991 International Conference on Parallel Processing*, pages II–46 – II–53, 1991.

[WT92] M. J. Wolfe and C. Tseng. The Power test for data dependence. *IEEE Transactions on Parallel and Distributed Systems*, 3(5):591–601, September 1992.