

Fractal Symbolic Analysis

Nikolay Mateev, Vijay Menon, Keshav Pingali
Department of Computer Science,
Cornell University, Ithaca, NY 14853

Abstract

Modern compilers perform wholesale restructuring of programs to improve their efficiency. *Dependence analysis* is the most widely used technique for proving the correctness of such transformations, but it suffers from the limitation that it considers only the memory locations read and written by a statement without assuming any particular interpretation for the operations in that statement. Exploiting the semantics of these operations permits more transformations to be proved correct, and is critical for automatic restructuring of codes such as LU with partial pivoting.

One approach to exploiting the semantics of program operations is *symbolic analysis and comparison* of programs. In principle, this technique is very powerful, but in practice, it is intractable for all but the simplest programs.

In this paper, we propose a new form of symbolic analysis and comparison of programs, appropriate for use in restructuring compilers. *Fractal* symbolic analysis compares a program and its transformed version by repeatedly simplifying these programs until symbolic analysis becomes tractable while ensuring that equality of the simplified programs is sufficient to guarantee equality of the original programs.

Fractal symbolic analysis combines some of the power of symbolic analysis with the tractability of dependence analysis. We discuss a prototype implementation of fractal symbolic analysis, and show how it can be used to solve the long-open problem of verifying the correctness of transformations required to improve the cache performance of LU factorization with partial pivoting.

1 Introduction

Modern compilers perform source-level transformations of programs to enhance locality and parallelism. Before a program is transformed, it must be analyzed to ensure that the proposed transformation does not change the meaning of

that program. Dependence analysis, the most commonly-used analysis technique, computes a partial order between execution instances of statements: a dependence is said to exist from one statement instance to another if one of these instances writes to a memory location that is read or written by the other one [23]. Any reordering of statements consistent with this partial order is permitted since it is guaranteed to leave the input-output behavior of the program unchanged.

Dependence analysis has been the focus of much research in the past two decades. In early work, program transformation was carried out manually by the programmer, and dependence analysis was used only to verify the legality of the transformation. More recently, the research community has invented powerful algorithms for automatically synthesizing performance-enhancing transformations for a program from representations of its dependences such as dependence matrices, cones, and polyhedra [3, 12, 14]. These techniques are sufficiently well-understood that they have been incorporated into production compilers such as the SGI MIPSPro.

In spite of these successes, dependence analysis has its shortcomings as is shown by the contrived program in Figure 1(a). In this program, there are dependences from statement S1 to S2, and from S1 to S3. There are only two statement reorderings consistent with this partial order: the original program, and the program obtained by reordering S2 and S3. In particular, the statement order shown in Figure 1(b) is not consistent with this partial order, so a compiler that relies on dependence analysis will declare that this transformation is illegal. However, it is not difficult to verify that the two programs in Figure 1 are equivalent. If x_{in} and y_{in} are the values of x and y at the start, the final values in x and y are $2 * x_{in}$ for both programs. Therefore, the statement reordering of Figure 1 is legal.

Intuitively, dependence analysis provides sufficient but not necessary conditions for the legality of restructuring transformations¹ since this analysis considers only the sets

⁰This work was supported by NSF grants CCR-9720211, EIA-9726388 and ACI-9870687.

¹A more precise variation of dependence analysis is *value-based* dependence analysis [8, 16, 18]. It is easy to verify that this alternative analysis

S1: $y = x$	$\Rightarrow$	S3: $x = 2 * x$
S2: $y = 2 * y$		S2: $y = 2 * y$
S3: $x = 2 * x$		S1: $y = x$

(a) Original program (b) Transformed program

Figure 1: Simple Reordering of Statements

of locations read and written by statements, and does not assume any particular interpretation (meaning) for the operations in the right-hand sides of these statements. Exploiting the semantics of these operations can lead to a richer space of program transformations as is shown by our simple example, and is critical for restructuring important codes like LU with partial pivoting.

Symbolic analysis is the usual way of exploiting the semantics of operations. To compare two programs for equality², we derive expressions for the outputs of these programs as functions of their inputs, and attempt to prove that these expressions are equal. If an algebraic axiom such as the distributive, associative, or commutative law of arithmetic can be assumed by the compiler, it may be used when proving equality.

In principle, symbolic analysis and comparison of programs is an extremely powerful technique for proving equality of programs; not only can be it used to verify the legality of program restructuring, but it can also be used to prove equality of programs that implement very different algorithms, such as sorting programs that implement quicksort and mergesort. However, for all but the simplest programs, symbolic execution and comparison is intractable.

In this paper, we describe *fractal* symbolic analysis which is a novel way of performing symbolic analysis and comparison of a program and its restructured version. Figure 2 illustrates the high-level idea. We assume we are given a symbolic analyzer that can symbolically analyze programs that are “simple enough”. Depending on the power of the analyzer, these may be programs with only straight-line code, or programs with only straight-line code and DO-ALL loops, etc. Let S be a program that is to be restructured to a program T . If these programs are simple enough, we invoke the symbolic analyzer on these programs and either prove or disprove their equality. On the other hand, if the programs are too complex to be analyzed by the symbolic analyzer, we generate two simplified programs S_1 and T_1 which have a very special property: *if these simplified programs are equal, we can conclude that the original programs S and T are equal as well*³. Analysis of the simplified programs is performed in a manner similar to the analysis of the original programs: if the programs are simple enough, they are analyzed by the symbolic analyzer; otherwise, they in turn are

makes no difference in our problem.

²In this paper, we are only interested in extensional equality: two programs are equal if their input-output behavior is equal.

³A minor technical detail is that a simplification step may actually produce a number of simplified programs from each of the original programs, in which case it is necessary to establish equality of that number of pairs of simplified programs.




Figure 2: Overview of Fractal Symbolic Analysis

simplified to produce new programs and so on (this is why we call our approach *fractal* symbolic analysis).

It is guaranteed that at some point, we will end up with programs S_n and T_n that are simple enough to be analyzed even by a symbolic analyzer that can only handle straight-line code. If we can prove that these programs are equal, we can conclude that S and T are equal; otherwise, we conservatively assume that S and T are not equal, and disallow the transformation.

There are two important caveats. First, the rules for simplifying programs are derived from the transformation that relates the two programs whose equality is to be established. Therefore, our approach cannot be used to prove that quicksort and mergesort are equal, for example, since these programs are not related by a restructuring transformation. Second, equality of simplified programs is a sufficient but not in general necessary condition for equality of the original programs. Therefore, successive simplification steps produce programs that are less and less likely to be equal even if the original programs are equal. It is desirable therefore that the core symbolic analyzer be powerful so that recursive simplification can be applied sparingly.

The rest of this paper is organized as follows. In Section 2, we introduce the highlights of our technology by discussing transformation of a small program. In Section 3, we discuss the simplification rules for key transformations in the literature. In Section 4, we describe the base symbolic analyzer we use in our implementation. We apply this technology to automatic blocking of LU factorization with pivoting in Section 5 and show that we achieve performance comparable with that of the LAPACK library [2] on the SGI Octane. Finally, in Section 6, we discuss ongoing work.

2 A Small Example

In this section, we discuss a small example which is a distillation of LU with partial pivoting and which illustrates various aspects of fractal symbolic analysis.

2.1 Source and Transformed Programs

The source program of Figure 3(a) traverses an array A ; at the j^{th} iteration, it swaps elements $A(j)$ and $A(j+1)$, and updates all the elements from $A(j+1)$ through $A(N)$ using the new value in $A(j)$. This is a much simplified version of LU factorization with partial pivoting in which entire rows of a matrix are swapped and entire sub-matrices are updated at each step. In our discussion, meta-variables $B1$ and $B2$ will be used to refer to the swap and update statement blocks respectively.

Loop distribution transforms this program into the one shown in Figure 3(b). In this program, all the swaps are done first, and then all the updates are done together. This transformation is useful because the second loop nest is perfectly-nested and can be tiled to get good locality of reference. Are these programs equal?




Figure 3: Loop Distribution in Running Example

Dependence analysis requires that there not be a dependence from an instance $B2(j2)$ to an instance $B1(j1)$ where $j1 > j2$. Unfortunately, this condition is violated: for any j_0 between 1 and $(N - 2)$, instance $B2(j_0)$ writes to location $A(j_0+1)$, and instance $B1(j_0+1)$ reads from it. Symbolic analysis of these programs on the other hand is too difficult.

2.2 Fractal Symbolic Analysis of Example

The key to fractal symbolic analysis is to consider the transformation of the program in Figure 3(a) to the one in Figure 3(b) not as a one-step transformation but as an incremental process in which instances of $B1$ are scheduled before successively earlier instances of $B2$ as shown in Figure 3(c). At each step of the incremental process, we locate an instance of block $B2$ (say $B2(j2)$) which is executed just *before* an instance of block $B1$ (say $B1(j1)$ where $j1 > j2$), and reschedule it so that it is executed just *after* that instance. It is obvious that at the end of this process, loop distribution is complete, and that this distribution is legal if each step of the process is legal.

To verify the legality of each step of the incremental process, we must verify that the programs shown in Figures 3(d) and (e) are equal; in other words, that $B1(j1)$ and $B2(j2)$ *commute* (assuming that $N > j1 > j2 \geq 1$). Notice that (i) these programs are simpler than the ones in Figures 3(a) and (b) since each one has one less loop, and (ii) their equality implies equality of the original programs, as required by Figure 2.

The symbolic analyzer we use in our implementation can analyze programs with straight-line code and DO-ALL loops, so it can perform analysis and comparison of the programs in Figure 3(d) and (e) without any further simplification. Let A_{in} and A_{out} be the values in array A before and after the execution of the program in Figure 3(d). It is easy to see that A_{out} can be expressed in terms of A_{in} as a *guarded symbolic expression* (GSE for short), shown in Figure 4, consisting of a sequence of guards defining array regions and symbolic expressions for the values in the array in those regions.

$$A_{out}(k) = \begin{cases} 1 \leq k \leq j2 & \rightarrow A_{in}(k) \\ k = j1 & \rightarrow A_{in}(j1+1)/A_{in}(j2) \\ k = j1+1 & \rightarrow A_{in}(j1)/A_{in}(j2) \\ \text{else} & \rightarrow A_{in}(k)/A_{in}(j2) \end{cases}$$

Figure 4: Guarded Symbolic Expression for A_{out}

An identical GSE expresses the result of executing the program of Figure 3(e). If A is assumed to be the only live variable after execution of the two programs, we can conclude that the programs of Figure 3(d) and (e) are equal, so the programs of Figure 3(a) and (b) are equal too.

2.3 Discussion

It is useful to understand what happens if we apply another step of simplification to the programs of Figure 3(d) and (e). The reordering of block $B1(j1)$ over the iterations of loop $B2(j2)$ can be viewed incrementally as a process in which

<pre> B1(j1)://swap tmp = A(j1); A(j1) = A(j1+1); A(j1+1) = tmp; B2(j2,i)://update body A(i) = A(i)/A(j2); </pre>	<pre> B2(j2,i)://update body A(i) = A(i)/A(j2); B1(j1)://swap tmp = A(j1); A(j1) = A(j1+1); A(j1+1) = tmp; </pre>
(a) $B1(j1); B2(j2,i)$	(b) $B2(j2,i); B1(j1)$

Figure 5: Another Step of Simplification

block $B1(j1)$ is moved over the iterations of loop $B2(j2)$ one iteration at a time. If each move is legal, the entire reordering is clearly legal. The simplified programs are shown in Figure 5(a) and (b); we must verify that $B1(j1)$ and $B2(j2,i)$ commute, assuming that $N > j1 > j2 \geq 1$ and that $N \geq i > j2$.

However, it is easy to verify that these programs are not equal. For example, for $k = i = j1$, the final values in $A_{out}(k)$ are different. This illustrates the caveat discussed in Section 1. Equality of the simplified programs is a sufficient but not in general necessary condition for equality of the original programs, so the simplification process that is at the heart of fractal symbolic analysis must be applied sparingly. In particular, had our base symbolic analyzer been able to analyze only straight-line code, we would have concluded conservatively that the loop distribution of Figure 3(a,b) is not legal.

To formalize the intuitive ideas presented in this section, we need to answer two questions.

1. What are the simplification rules?
2. What core symbolic analyzer is both powerful and practical?

We answer these questions next.

3 Simplification Rules

To determine equality of a program and its restructured version, the compiler uses a table similar to the one shown in Table 1(a). To verify legality of a transformation shown in the first column, the Commute function is invoked on appropriate instantiations of program statements as shown in the second column of this table. The intuition behind these legality conditions has been described in Section 2 — the transformation is viewed incrementally as reordering pairs of instances at a time. For lack of space, we have shown the legality conditions for only a few common transformations.

The Commute function is shown in Figure 6. It is passed two program fragments pgm_1 and pgm_2 , some optional *bindings* which are constraints on free variables in the program fragments, and a list of variables that are live at the end of execution of the program fragments. It returns true if pgm_1 and pgm_2 commute — that is, if the result of executing pgm_1 followed by pgm_2 is the same as the result of executing them in the opposite order, assuming that we are

only interested in the values of live variables at the end of execution.

The Commute function invokes the symbolic analysis and comparison engine directly by calling the Compare method if the program fragments pgm_1 and pgm_2 are simple enough. Our implementation of this method is discussed in Section 4. On the other hand, if these program fragments are not simple enough, it applies the rules shown in Table 1(b) to simplify the programs recursively till the programs are simple enough to be analyzed symbolically. As with the rules shown in Table 1(a), the rules in Table 1(b) are easy to understand if the restructuring is viewed incrementally. Note that exactly one of the rules in Table 1(b) can be applied at each simplification step, resulting in a deterministic simplification procedure.

3.1 Proof of Correctness of Simplification Rules

The validity of the legality conditions in Table 1 follows from the following result, many variations of which have appeared in the literature [11].

Lemma 1 *Let $S = \{S_1; S_2; S_3; \dots; S_n\}$ be a sequence of program fragments, and let p be any permutation on S . Define $R(p) = \{(S_i, S_j) : 1 \leq i < j \leq n \wedge p(i) > p(j)\}$ as the set of pairs of statements reordered by p . Then, the program S is equivalent to the program $p(S) = \{S_{p(1)}; S_{p(2)}; S_{p(3)}; \dots; S_{p(n)}\}$ if $\{S_i; S_j\}$ is equivalent to $\{S_j; S_i\}$ where $(S_i, S_j) \in R(p)$.*

Proof: The Lemma follows from induction on the cardinality of $R(p)$. If $\|R(p)\| = 0$ then clearly $S = p(S)$. Otherwise, there must exist an i such that $p(i+1) < p(i)$, which implies that $(S_{p(i+1)}, S_{p(i)}) \in R(p)$. Since $S_{p(i)}$ and $S_{p(i+1)}$ commute and are adjacent in $p(S)$, transposing them gives us an equivalent program $p'(S)$ where $R(p') \subset R(p)$ and $\|R(p')\| = \|R(p)\| - 1$. By induction, $S = p(S)$. $\square$

Intuitively, Lemma 1 allows us to reformulate the problem of checking legality of a transformation as a problem of verifying commutativity of *statement instances* that are reordered by that transformation. The validity of the rules given in Table 1 follows directly from this result.

3.2 Reduction Example

An interesting example of the use of the rules in Table 1 is provided by the reduction example shown in Figure 7. The loop interchange shown in Figures 7(a,b) is legal if we assume that addition is commutative and associative, but every iteration of the loop nest reads and writes variable k , so dependences are violated by the restructuring. This program is simple enough that symbolic analysis is tractable in principle although it requires reasoning about summations. In

Transformation	Legality Condition
Statement Reordering $S1; S2; \quad \leftrightarrow \quad S2; S1;$	$Commute(\langle S1, S2 \rangle)$
Loop Distribution/Jamming $\text{do } i = 1, n$ $\quad S1(i); \quad \leftrightarrow \quad \text{do } i = 1, n$ $\quad S1(i); \quad \text{do } i = 1, n$ $\quad S2(i);$	$Commute(\langle S1(i_1), S2(i_2) \rangle : 1 \leq i_2 < i_1 \leq n)$
Loop Interchange $\text{do } i = 1, n$ $\quad \text{do } j = 1, m$ $\quad \quad S(i, j); \quad \leftrightarrow \quad \text{do } j = 1, m$ $\quad \quad \text{do } i = 1, n$ $\quad \quad \quad S(i, j);$	$Commute(\langle S(i_1, j_1), S(i_2, j_2) \rangle : 1 \leq i_1 < i_2 \leq n \wedge 1 \leq j_2 < j_1 \leq m)$
Linear Loop Transformation T $\text{do } (i1, i2, \dots, ik)$ $\quad S(i1, i2, \dots, ik); \quad \leftrightarrow \quad \text{do } (i1', i2', \dots, ik')$ $\quad S'(i1', i2', \dots, ik'); \quad \text{where } (i1', i2', \dots, ik') = T(i1, i2, \dots, ik)$	$Commute(\langle S(\vec{i}), S(\vec{j}) \rangle : \vec{i} \prec \vec{j} \wedge T(\vec{i}) \succ T(\vec{j}))$

(a) Simplification Rules for Common Loop Transformations

Commute Condition	Recursive Condition
Statement Sequence $Commute(\langle S1; S2; \dots; SN, B2 \rangle : cond)$	$Commute(\langle S1, B2 \rangle : cond) \wedge$ $Commute(\langle S2, B2 \rangle : cond) \wedge$ $\dots$ $Commute(\langle SN, B2 \rangle : cond)$
Loop $Commute(\langle \text{do } i = 1, u$ $\quad \quad S1(i); \quad , B2 \rangle : cond)$	$Commute(\langle S1(i), B2 \rangle : cond \wedge 1 \leq i \leq u)$
Conditional Statement $Commute(\langle \text{if } (pred) \text{ then}$ $\quad \quad S1;$ $\quad \text{else}$ $\quad \quad S2; \quad , B2 \rangle : cond)$	$Commute(\langle S1, B2 \rangle : cond \wedge pred) \wedge$ $Commute(\langle S2, B2 \rangle : cond \wedge \neg pred)$

(b) Recursive Simplification Rules

Table 1: Simplification Rules

practice most compilers use pattern matching to detect reductions and treat them specially, but pattern matching is notoriously fragile (for example, introducing a temporary into the reduction will cause most pattern matchers to fail).

Fractal symbolic analysis solves the problem elegantly. Using Table 1, we generate the simplified programs shown in Figure 7(c,d). A symbolic analyzer that can analyze straight-line code can deduce easily that these programs equal if it is allowed to assume that addition is commutative and associative; if addition cannot be assumed to have these properties, the symbolic analysis and comparison engine deduces correctly that the program transformation is not legal.

3.3 Discussion

As mentioned in Section 2, the precision of fractal symbolic analysis depends on the power of the core symbolic comparison engine. Notice that the procedure in Figure 6 stops simplifying as soon as the statements being compared can be handled by the Compare procedure. A more powerful Compare procedure will result in fewer levels of simplification and potentially more accurate symbolic analysis.

4 Symbolic Analysis and Comparison Engine

We now describe the core symbolic comparison procedure that is invoked after simplification. As mentioned earlier, it is important for this procedure to be as powerful as is tractable so that simplification can be applied sparingly. In our work, we have found that it is sufficient if the symbolic comparison procedure can handle the following class of programs⁴.

- Programs consist of assignment statements, for-loops and conditionals. No unstructured control flow is allowed.
- Loops do not have loop-carried dependences.
- Array indices and loop bounds are restricted to be affine functions of enclosing loop variables and symbolic constants, and predicates are restricted to be conjunctions and disjunctions of affine inequalities.

⁴It is important to remember that these restrictions are only for programs that can be analyzed directly by the symbolic analyzer; they are not required for input programs to fractal symbolic analysis since complex programs are simplified recursively till they meet the restrictions for the core symbolic engine.

```

Commute(pgm1,pgm2,bindings,livevars) {
  if Simple(pgm1) then
    if Simple(pgm2)
      return Compare({pgm1;pgm2}, {pgm2;pgm1},
        bindings,livevars)
    else
      return Commute(pgm2,pgm1,bindings,livevars)
  else
    case (pgm1) {
      ⟨pgm'1;pgm''1⟩ →
        return Commute(pgm'1,pgm2,
          bindings,livevars) ∧ Commute(pgm''1,pgm2,
            bindings,livevars)

      ⟨if pred then pgm'1 else pgm''1⟩ →
        return Commute(pgm'1,pgm2,
          bindings||pred,livevars) ∧ Commute(pgm''1,
            pgm2,bindings||¬pred,livevars)

      ⟨do i = l, u pgm'1(i)⟩ →
        return Commute(pgm'1(i),pgm2,
          bindings||∀i.l ≤ i ≤ u,livevars)
    }
}

```

Figure 6: The *Commute* function

<pre> do i = 1, M do j = 1, N B(i, j) ::= update on k k = k + A(i, j) (a) Original Program </pre>	<pre> do j = 1, N do i = 1, M B(i, j) ::= update on k k = k + A(i, j) (b) Transformed Program </pre>
---	--

<pre> B(i1, j1): k = k + A(i1, j1) B(i2, j2): k = k + A(i2, j2) (c) B(i1, j1); B(i2, j2) </pre>	<pre> B(i2, j2): k = k + A(i2, j2) B(i1, j1): k = k + A(i1, j1) (d) B(i2, j2); B(i1, j1) </pre>
---	---

Figure 7: Loop Interchange of Reduction Code

$$A(\vec{k}) = \begin{cases} guard_1(\vec{k}) \rightarrow expression_1(\vec{k}) \\ guard_2(\vec{k}) \rightarrow expression_2(\vec{k}) \\ \vdots \\ guard_n(\vec{k}) \rightarrow expression_n(\vec{k}) \end{cases}$$

Figure 8: Guarded Symbolic Expressions

The important constraint is the second one. Although a loop may write to a section of an array that is potentially unbounded at compile-time, at most one iteration may affect the value of any given location in an array. This ensures that the symbolic value of a given element of the array can be expressed finitely. Techniques such as scalar expansion [23] should be used to aggressively eliminate loop-carried dependencies.

We can then summarize the unbounded set of expressions for the values in an entire array with a finite expression called a *guarded symbolic expression* (or GSE for short) which contains symbolic expressions that hold for affinely constrained portions of the array as shown in Figure 8. Figure 4 shows an example of a GSE.

```

Compare(stmt1,stmt2,bindings,livevars) {
  live1 = set of live altered variables in stmt1
  live2 = set of live altered variables in stmt2
  if(live1 ≠ live2)
    return false

  for each a( $\vec{k}$ ) in live1 {
    tree1 = Build_Expr_Tree(stmt1,a( $\vec{k}$ ),∅)
    tree2 = Build_Expr_Tree(stmt2,a( $\vec{k}$ ),∅)

    gse1 = Normalize_GSE(tree1,bindings)
    gse2 = Normalize_GSE(tree2,bindings)

    if(¬ Compare_GSEs(gse1,gse2)
      return false
  }
  return true
}

```

Figure 9: Comparison of Simple Programs

```

Build_Expr_Tree(stmt,tree,bindings) {
  case (tree) {
    Op(op,tree1,...,treen) →
      return Op(op,
        Build_Expr_Tree(stmt,tree1,bindings),...,
        Build_Expr_Tree(stmt,treen,bindings))

    Cond(pred,tree1,tree2) →
      return Cond(pred,
        Build_Expr_Tree(stmt,tree1,bindings),
        Build_Expr_Tree(stmt,tree2,bindings))

    A( $\vec{k}$ ) →
      case (stmt) {
        ⟨A'(T ·  $\vec{i}$  + c) = tree1( $\vec{i}$ )⟩ →
          if (A = A') then
            return Cond(bindings|| $\vec{k} = T \cdot \vec{i} + c$ ,
              tree1(T-1 · ( $\vec{k} - c$ )),A( $\vec{k}$ ))
          else
            return A( $\vec{k}$ )

        ⟨stmt1;stmt2⟩ →
          return Build_Expr_Tree(stmt1,
            Build_Expr_Tree(stmt2,tree,bindings),
            bindings)


        ⟨if pred then stmt1 else stmt2⟩ →
          return Cond(bindings||pred,
            Build_Expr_Tree(stmt1,tree,bindings),
            Build_Expr_Tree(stmt2,tree,bindings))

        ⟨do ik = lk, uk stmt1⟩ →
          return Build_Expr_Tree(stmt1,tree,
            bindings||∃ik.lk ≤ ik ≤ uk)
      }
  }
}

```

Figure 10: Expression Tree Generation

Figure 9 provides a high-level overview of our symbolic comparison algorithm. We consider each altered scalar or array variable in the two programs being compared. Note that we only need to consider *live* variables [1]. If the GSE's corresponding to each live altered variable are equal, the two programs are declared to be equal. We now describe how GSE's are constructed and compared.



(a) Tree for Figure 3(d) (b) Tree for Figure 3(e)

Figure 11: Conditional Expression Trees for Running Example

4.1 Generation of Conditional Expression Trees

A guarded symbolic expression is essentially a description of the effect of a program on an array. As an intermediate step towards the construction of this description, we build a symbolic representation of the program that we call a *conditional expression tree*. A conditional expression tree may be viewed as a functional expression that describes the output of the program as a function of its inputs. Figure 11 shows the conditional expression trees for the programs of Figure 3(d,e). The interior nodes of the tree are predicates and operators while the leaves of the tree are scalars and array references. Since simplified programs have only straight-line code and DO-ALL loops, the construction of these trees is straight-forward. Figure 10 shows an algorithm to generate such trees. This algorithm processes the statements of a program in reverse order, determining at each step the tree corresponding to relevant output data in terms of input data and linking these together to produce the final result.

4.2 Normalization to Guarded Symbolic Expressions

In general, the conditional expression trees generated above contain a mix of conditions predicated by affine constraints on one hand and arithmetic expressions on the other (Figure 11(a) is an example of such a mixture). To convert these to guarded symbolic expressions, we need to separate the two (Figure 11(b) is an example where the affine constraints are separate from the arithmetic expressions). We accomplish this by factoring the affine constraints outside of the symbolic operations by repeated application of the following transformation.

$$\begin{aligned} & \text{Op}(op, tree_1, \dots, \text{Cond}(pred, tree_{t_i}, tree_{f_i}), \dots, tree_n) \\ & \quad \downarrow \\ & \text{Cond}(pred, \text{Op}(op, tree_1, \dots, tree_{t_i}, \dots, tree_n), \\ & \quad \text{Op}(op, tree_1, \dots, tree_{f_i}, \dots, tree_n)) \end{aligned}$$

At this point, the guards are generated by combining the predicates at the top of the factored expression tree, and the corresponding arithmetic expressions are simply taken from the subtrees beneath these predicates, as is shown in Fig-

```
Factor(tree) {
  case (tree) {
    Op(op, tree_1, ..., Cond(pred, tree_{t_i}, tree_{f_i}), ..., tree_n) →
      return Factor(Cond(pred,
        Op(op, tree_1, ..., tree_{t_i}, ..., tree_n),
        Op(op, tree_1, ..., tree_{f_i}, ..., tree_n)))

    Op(op, tree_1, ..., tree_n) →
      return Op(op, Factor(tree_1), ..., Factor(tree_n))

    Cond(pred, tree_i, tree_f) →
      return Cond(pred, Factor(tree_i), Factor(tree_f))

    A(k) →
      return A(k)
  }
}

Build_GSE(tree, guard)
if (guard) then
  case (tree) {
    Cond(pred, tree_t, tree_f) →
      return Build_GSE(tree_t, guard ∧ pred) ∪
        Build_GSE(tree_f, guard ∧ ¬pred)

    Op(op, tree_1, ..., tree_n) →
      A(k) →
        return {(guard, expr)}
  }
else
  return ∅
}

Normalize_GSE(tree, bindings) {
  return Build_GSE(Factor(tree), bindings)
}
```

Figure 12: From Expression Trees to GSE's

```
Compare_GSEs(gse_1, gse_2) {
  for each (guard_1, expr_1) in gse_1 {
    for each (guard_2, expr_2) in gse_2 {
      if (guard_1 ∧ guard_2 ≠ false)
        if (expr_1 ≠ expr_2)
          return false
    }
  }
  return true
}
```

Figure 13: Comparison of GSE's

ure 12.

4.3 Comparison of Guarded Symbolic Expressions

Finally, Figure 13 illustrates the comparison of two guarded symbolic expressions. There are two steps to this comparison. First, we must compare each pair of affine guards of the two guarded symbolic expressions. Second, for any two guards that potentially intersect, we must compare the corresponding symbolic expressions. If all such symbolic expression match, then the guarded symbolic expressions are declared to be equal. The validity of this conclusion follows from the following argument. Each guard specifies some region of the index space of the array in question, and the union of these regions in a guarded symbolic expression is

equal to the entire index space of that array. If the values in the two guarded symbolic expressions are identical whenever their guards intersect, the two array values are obviously equal.

For comparison of affine guards, we may employ an integer programming tool such as the Omega Library [17]. If the tool proves that a pair of affine guards do not intersect, no comparison of the corresponding arithmetic expressions needs to be performed. On the other hand, if the guards so intersect, the expressions must be compared for equality. In our current implementation, we just check for syntactic equality. This is sufficient both for our simple example and, as we shall see in Section 5, for LU factorization. It would be easy to use Maple or some other symbolic algebra tool for this test if we wished to exploit algebraic properties of numbers.

5 LU with Pivoting

Fractal symbolic analysis was developed for use in an ongoing project on optimizing the cache behavior of dense numerical linear algebra programs. LU factorization with partial pivoting is a key routine in this application area since it is used to solve systems of linear equations of the form $Ax = b$. Figure 14 shows the canonical version of LU factorization with pivoting that appears in the literature [9]. In iteration j of the outer loop, computations are performed on column j of the matrix A , and a portion of the matrix to the right of this column is updated.

LU factorization with pivoting poses a number of challenges for compiler writers.

1. Cache-optimized versions of LU factorization can be found in the LAPACK library [2]. These *blocked codes* are too complex to be reproduced here, but they perform much better than the *point version* shown in Figure 14.

Given point-wise LU factorization with pivoting, can a compiler automatically generate a cache-optimized version by blocking the code? If so, how does the performance of the compiler-optimized code compare with that of hand-blocked code?

2. Right-looking (eager updates) and left-looking (lazy updates) versions of LU factorization can be obtained by interchanging the two loops of the update step. Can a compiler transform right-looking LU to left-looking LU and vice versa?

Fractal symbolic analysis is crucial to address both these challenges. For lack of space, we discuss only the problem of blocking.

```

do j = 1, N
  // Pick the pivot
  p(j) = j
  do i = j+1, N
    if abs(A(i,j)) > abs(A(p(j),j))
      p(j) = i
  // Swap rows
  do k = 1, N
    tmp = A(j,k)
    A(j,k) = A(p(j),k)
    A(p(j),k) = tmp
  // Scale current column
  do i = j+1, N
    A(i,j) = A(i,j) / A(j,j)
  // Update portion of matrix to right of column j
  do k = j+1, N
    do i = j+1, N
      A(i,k) = A(i,k) - A(i,j)*A(j,k)

```

Figure 14: LU Factorization with Pivoting

5.1 Automatic Blocking of LU Factorization

To obtain code competitive with LAPACK code, Carr and Lehoucq suggest carrying out the following sequence of restructuring transformations [4].

1. Stripmine the outer loop to formulate block-column operations.
2. Index-set-split the expensive update operation to separate computation outside the current block-column from computation inside the current block-column.
3. Distribute the inner of the stripmined loops to isolate the out-of-column update.
4. Tile the out-of-column update.

The first two steps, stripmining and index-set-splitting, are trivially legal as they do not reorder any computation. The next step, loop distribution, is not necessarily legal. If legality is checked using dependence analysis, the compiler will declare the distribution illegal if there is a dependence from an iteration $B2(m)$ to an iteration $B1(l)$ where $l > m$. In fact, such a dependence exists in our program; for example, both $B2(j)$ and $B1(j+1)$ read and write to $A(m+1, jB+B \dots N)$. *Therefore, a compiler that relies on dependence analysis cannot block LU with pivoting using the transformation strategy of Carr and Lehoucq.* Carr and Lehoucq suggest that a compiler may be endowed with application-specific information to recognize the swap and update operations in LU factorization, and to realize that they can be legally interchanged.

We now discuss how our implementation of fractal symbolic analysis, a general-purpose technique, can verify the legality of these transformations.

To verify the legality of the loop distribution step, our compiler consults Table 1(a) and determines that it must check if $B1(l)$ commutes with $B2(m)$ where $jB \leq m < l \leq jB + B - 1$, as shown in Figure 16. The compiler invokes the Commute method in Figure 6 with these parameters, as shown in Figure 17. However, these simpler programs are not “simple enough”; the loop that computes the pivot in $B1.b(1)$ is a recurrence that cannot be handled by our core symbolic comparison engine, as we discussed

```

do jB = 1, N, B
do j = jB, jB+B-1
B1(j):
  // Pick the pivot
  p(j) = j
  do i = j+1, N
    if abs(A(i,j)) > abs(A(p(j),j))
      p(j) = i
  // Swap rows
  do k = 1, N
    tmp = A(j,k)
    A(j,k) = A(p(j),k)
    A(p(j),k) = tmp
  // Scale column
  do i = j+1, N
    A(i,j) = A(i,j) / A(j,j)
  // In-Column Update
  do k = j+1, jB+B-1
    do i = j+1, N
      A(i,k) = A(i,k) - A(i,j)*A(j,k)
B2(j):
  // Right-Looking Update
  do k = jB+B, N
    do i = j+1, N
      A(i,k) = A(i,k) - A(i,j)*A(j,k)

```

(a) Before Loop Distribution

```

do jB = 1, N, B
do j = jB, jB+B-1
B1(j):
  // Pick the pivot
  p(j) = j
  do i = j+1, N
    if abs(A(i,j)) > abs(A(p(j),j))
      p(j) = i
  // Swap rows
  do k = 1, N
    tmp = A(j,k)
    A(j,k) = A(p(j),k)
    A(p(j),k) = tmp
  // Scale column
  do i = j+1, N
    A(i,j) = A(i,j) / A(j,j)
  // In-Column Update
  do k = j+1, jB+B-1
    do i = j+1, N
      A(i,k) = A(i,k) - A(i,j)*A(j,k)
  // Distributed Loop
  do j = jB, jB+B-1
B2(j):
  // Right-Looking Update
  do k = jB+B-1, N
    do i = j+1, N
      A(i,k) = A(i,k) - A(i,j)*A(j,k)

```

(b) After Loop Distribution

Figure 15: LU Factorization: Distribution Step

```

B1.a(1): p(1) = 1
B1.b(1): do i = 1+1, N
  if abs(A(i,1)) > abs(A(p(1),1))
    p(1) = i
B1.c(1): do k = 1, N
  tmp = A(1,k)
  A(1,k) = A(p(1),k)
  A(p(1),k) = tmp
B1.d(1): do i = 1+1, N
  A(i,1) = A(i,1) / A(1,1)
B1.e(1): do k = 1+1, jB+B-1
  do i = 1+1, N
    A(i,k) = A(i,k) - A(i,1)*A(1,k)
B2(m): do k = jB+B, N
  do i = m+1, N
    A(i,k) = A(i,m) - A(i,m)*A(m,k)

```

(a) B1(1); B2(m)

```

B2(m): do k = jB+B, N
  do i = m+1, N
    A(i,k) = A(i,m) - A(i,m)*A(m,k)
B1.a(1): p(1) = 1
B1.b(1): do i = 1+1, N
  if abs(A(i,1)) > abs(A(p(1),1))
    p(1) = i
B1.c(1): do k = 1, N
  tmp = A(1,k)
  A(1,k) = A(p(1),k)
  A(p(1),k) = tmp
B1.d(1): do i = 1+1, N
  A(i,1) = A(i,1) / A(1,1)
B1.e(1): do k = 1+1, jB+B-1
  do i = 1+1, N
    A(i,k) = A(i,k) - A(i,1)*A(1,k)

```

(b) B2(m); B1(1)

Figure 16: Simplified Comparison #1

in Section 4. Therefore, these programs are simplified again using the rule for statement sequences in Table 1(b). This requires the compiler to test whether $B2(m)$ commutes with the five subblocks in $B1(1)$. Other than $B1.c(1)$, the subblocks of $B1(1)$ touch data that is disjoint from the data touched by $B2(m)$. Therefore, our compiler deduces that these subblocks commute with $B2(m)$ (a small detail is that the analysis of whether $B1.b(1)$ commutes with $B2(m)$ requires an additional step of simplification to eliminate the recurrence in $B1.b(1)$).

The only problem remaining is to demonstrate that $B1.c(1)$ and $B2(m)$ commute as shown in Figure 18. At this point, these programs are “simple enough”, and the Compare method in Figure 9 is invoked to establish equality of the simplified programs. In fact, they are quite similar to the running example of Section 2. The only live, altered variable in either program is the array A , and the Compare

method generates guarded symbolic expressions for A from each program. Both GSE’s generated from Figure 18 contain six guarded regions, shown pictorially in Figure 19. To prove that the GSE’s are actually equivalent, Compare_GSEs is invoked to generate the 36 pairwise intersections, and the Omega library [17] is used to test non-emptiness of these regions. Only six intersections are non-empty (the six regions shown in Figure 19), and the corresponding symbolic expressions are syntactically identical in each case. Thus, the compiler is able to demonstrate the equality of the simplified programs and, therefore, the equality of the programs in Figure 15. Since the symbolic expressions are syntactically equal, it follows that the restructuring does not change the output of the program even if arithmetic is finite-precision.

One important note is that the programs of Figure 18 are equal only if $p(j) \geq j$. Techniques such as value propagation [7, 15] have been developed to perform this type of anal-




Figure 17: Commute Calls for LU

```

B1.c(l): do k = 1, N
          tmp = A(1, k)
          A(1, k) = A(p(1), k)
          A(p(1), k) = tmp
B2(m):   do k = jB+B, N
          do i = m+1, N
            A(i, k) = A(i, m) - A(i, m)*A(m, k)

```

(a) B1.c(l); B2(m)

```

B2(m):   do k = jB+B, N
          do i = m+1, N
            A(i, k) = A(i, m) - A(i, m)*A(m, k)
B1.c(l): do k = 1, N
          tmp = A(1, k)
          A(1, k) = A(p(1), k)
          A(p(1), k) = tmp

```

(b) B2(m); B1.c(l)

Figure 18: Simplified Comparison #2




Figure 19: Regions and Expressions for Simplified LU

ysis for indirect array accesses to more accurately compute dependences. It is clear that this information may easily be inferred from the pivot computation in B1.a and B1.b. In our implementation, this information is passed by the compiler as bindings to the method Commute along with the legality conditions in Table 1.

With this information, our implementation of fractal symbolic analysis is able to automatically establish the legality of the loop distribution transformation in Figure 15. For this example, our implementation, prototyped in Caml-Light [13], took slightly less than one second. Most of the analysis time is spent on the construction and comparison of guarded symbolic expressions since we have not yet optimized the code.

5.2 Experimental Results

To study the effects of automatic blocking, we ran the SGI compiler on the LU code in Figure 15, after loop distribution is performed. Given just this slightly transformed code, the SGI compiler is now able to produce significantly faster code that does not degrade as matrices exceed the cache. Once the loop is distributed, the compiler is able to automatically tile the right-looking update (B2) and essentially accomplish the last Carr/Lehoucq step listed above.

Once we have isolated the update, dependence analysis will allow us to distribute the portion of the swap computation to left and right of the current block column outside the block column computation to obtain a blocking structure almost identical to that in the Netlib LAPACK. The SGI compiler does not do this itself; when we apply this transformation by hand, we see a modest increase in performance.

Nevertheless, this code, at 200 MFlops, is still a factor of two slower than the LAPACK codes. Further experimentation found the remaining performance gap due the compiler's suboptimal treatment of the right-looking update computation. Although, the SGI compiler is able to now block the update, we surmise that it may have been confused by the partially triangular loop bounds of the update. When we index-set split the i loop by hand to separate the triangular and rectangular portions of the update, the compiler generated substantially faster code achieving over 300 MFlops. Finally, we note that if we replace the triangular and rectangular portions of the update with the corresponding BLAS-3 calls (DTRSM and DGEMM) used in LAPACK, the resulting code achieves nearly 400 MFlops and is within 10% of




Figure 20: Transforming LU with Pivoting

Netlib LAPACK and 20% of the best hand-optimized code. Thus, with the ability to isolate the update, we believe that compilers should be able to nearly match the LAPACK as their ability to match the performance of BLAS on perfectly nested codes improves.

6 Related and Future Work

A simple kind of symbolic analysis called *value numbering* [1] and a generalization called *global value numbering* [20] are used in some optimizing compilers to identify opportunities for common subexpression elimination and constant propagation, but these techniques are not useful for comparing *different* programs.

Sophisticated symbolic analysis techniques for finding *generalized induction variables* have been developed by Haghighat and Polychronopoulos [10] and by Rauchwerger and Padua [19], but their goal is to perform strength increasing to eliminate loop-carried dependences for automatic parallelization. Since this may produce DO-ALL loops from loops with loop-carried dependences, it may be advantageous to preprocess programs in this way before applying fractal symbolic analysis since this eliminates the need for recursive simplification in such programs.

The work that is closest to our own is Rinard’s *commutativity analysis* [21] which is a program parallelization technique that uses symbolic analysis to determine if method invocations can be executed concurrently. This approach is based on the insight that a sequence of atomic operations can be executed in parallel if each pair of operations can be shown to commute. We are interested in proving the correctness of program transformations, not in parallelizing programs, and requiring all operations to commute with each other is too strong a condition for our application; there is also no analog of recursive simplification in commutativity analysis.

The algorithm for generating guarded symbolic expressions in Section 4 is reminiscent of backward slicing [22]

which is a technique that isolates the portion of a program that may affect the value of a variable at some point in the program. Our algorithm is simpler than the usual algorithms for backward slicing since the programs it must deal with have been simplified beforehand by recursive simplification, an operation that has no analog in backward slicing. A similar statement can be made about the computation of *last-write* information [6].

The work described in this paper can be extended in many ways such as the following.

- The symbolic analysis engine can be extended to recognize and summarize reductions involving associative arithmetic operations like addition and multiplication, and the symbolic comparison engine can invoke a symbolic algebra tool like Maple [5] to compare such expressions using the usual laws of arithmetic. These enhancements might eliminate the need for recursive simplification in some programs, but we do not yet have any applications where this additional power is needed.
- The intuition behind fractal symbolic analysis is to view a program transformation as a process which transforms the initial program incrementally to the final program. In general, there are many incremental processes that achieve the effect of a given transformation, and Table 1 shows just one such way for the transformations listed there. Given a program and its transformation, fractal symbolic analysis may succeed in proving the correctness of some of these processes, but it may fail conservatively for others. A useful analogy is induction which can be viewed as an incremental way of proving predicates: given a predicate, we can usually formulate many inductive strategies for proving it, but only some of them will actually succeed. Is it useful to explore an entire space of incremental processes for converting one program to another? If so, how do we manage the search to keep it tractable?
- The proof of correctness of the transformation for LU with pivoting discussed in Section 5 required knowing that $p(j) \geq j$. This constraint is easy to deduce, but how does a compiler know in general that this information is useful? One approach is to have the compiler gather as many constraints on variables as it can deduce, and pass them to the fractal symbolic analyzer. An alternative lazy strategy is to gather only facts that are required for proving the validity of transformations, but it is not clear how such facts can be identified.
- Finally, we note that dependence information for loops can be represented abstractly using dependence vectors, cones, polyhedra etc. These representations have been exploited to *synthesize* transformation sequences [3, 12, 14]. At present, we do not know suitable representations for the results of fractal symbolic analysis, nor do

we know how to synthesize transformation sequences from such information.

References

- [1] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques and Tools*. Addison Wesley, Reading, MA, second edition, 1986.
- [2] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen, editors. *LAPACK Users' Guide. Second Edition*. SIAM, Philadelphia, 1995.
- [3] U. Banerjee. A theory of loop permutations. In *Languages and compilers for parallel computing*, pages 54–74, 1989.
- [4] S. Carr and R. B. Lehoucq. Compiler blockability of dense matrix factorizations. *ACM Transactions on Mathematical Software*, 23(3):336–361, September 1997.
- [5] B. Char, K. Geddes, and G. Gonnet. The Maple symbolic computation system. *SIGSAM Bulletin (ACM Special Interest Group on Symbolic and Algebraic Manipulation)*, 17(3/4):31–42, August/November 1983.
- [6] J.-F. Collard and M. Griebel. A precise fixpoint reaching definition analysis for arrays. In *Proc. of 12th International Workshop on Languages and Compilers for Parallel Computing, (LCPC99)*, August 1999.
- [7] L. A. DeRose. Compiler techniques for MATLAB programs. Technical Report 1956, Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois, May 1996.
- [8] P. Feautrier. Dataflow analysis of array and scalar references. *International Journal of Parallel Programming*, 20(1):23–53, February 1991.
- [9] G. Golub and C. V. Loan. *Matrix Computations*. The Johns Hopkins University Press, 1996.
- [10] M. R. Haghighat and C. D. Polychronopoulos. Symbolic analysis for parallelizing compilers. *ACM Transactions on Programming Languages and Systems*, 18(4):477–518, July 1996.
- [11] S. M. Johnson. Generation of permutations by adjacent transposition (in Technical Notes and Short Papers). *Mathematics of Computation*, 17(83):282–285, July 1963.
- [12] M. S. Lam, E. E. Rothberg, and M. E. Wolf. The cache performance and optimizations of blocked algorithms. In *Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 63–74, Apr. 8–11, 1991.
- [13] X. Leroy. The Caml Light system, release 0.71. Documentation and user's manual. Technical report, INRIA, March 1996. Available from *Projet Cristal* at <http://pauillac.inria.fr>.
- [14] W. Li and K. Pingali. A singular loop transformation based on non-singular matrices. *International Journal of Parallel Programming*, 22(2), April 1994.
- [15] V. Maslov. Enhancing array dataflow dependence analysis with on-demand global value propagation. In *Proc. International Conference on Supercomputing*, pages 265–269, July 1995.
- [16] D. E. Maydan, S. P. Amarasinghe, and M. S. Lam. Data dependence and data-flow analysis of arrays. In *5th Workshop on Languages and Compilers for Parallel Computing*, pages 434–448, August 1992.
- [17] W. Pugh. The Omega test: A fast and practical integer programming algorithm for dependence analysis. In *Communications of the ACM*, pages 102–114, Aug. 1992.
- [18] W. Pugh and D. Wonnacott. An exact method for analysis of value-based array data dependences. Technical Report CS-TR-3196, University of Maryland, College Park, Portland, Ore., December 1993.
- [19] L. Rauchwerger and D. A. Padua. The LRPD test: Speculative run-time parallelization of loops with privatization and reduction parallelization. *IEEE Transactions on Parallel and Distributed Systems*, 10(2), February 1999.
- [20] J. H. Reif and R. E. Tarjan. Symbolic program analysis in almost linear time. *SIAM Journal on Computing*, 11(1):81–93, February 1982.
- [21] M. C. Rinard and P. C. Diniz. Commutativity analysis: a new analysis technique for parallelizing compilers. *ACM Transactions on Programming Languages and Systems*, 19(6):942–991, November 1997.
- [22] M. Weiser. Program slicing. *IEEE Transactions on Software Engineering*, 10(4):352–357, 1984.
- [23] M. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley Publishing Company, 1995.