

Enhancing the Role of Inlining in Effective Interprocedural Parallelization

Jichi Guo Mike Stiles Qing Yi Kleanthis Psarris
 University of Texas at San Antonio
 {jguo,mstiles,qingyi,psarris}@cs.utsa.edu

Abstract—The emergence of multi-core architectures makes it essential for optimizing compilers to automatically extract parallelism for large scientific applications composed of many subroutines residing in different files. Inlining is a well-known technique which can be used to erase procedural boundaries and enable more aggressive loop parallelization. However, conventional inlining cannot be applied to external libraries where the source code is not available, and when overly applied, it can degrade the effectiveness of compiler optimizations due to excessive code complexity. This paper highlights some obstacles we encountered while applying conventional inlining combined with automatic loop parallelization using the Polaris optimizing compiler and presents a new approach, annotation-based inlining, to effectively overcome these obstacles. Our experimental results show that the annotation-based inlining approach can eliminate negative impact of conventional inlining while enhancing the effectiveness of interprocedural parallelization for a majority of applications from the PERFECT benchmark suite.

I. INTRODUCTION

As multi-core architectures become ubiquitous in modern computing, optimizing compilers need to automatically extract parallelism for large scale scientific applications composed of many subroutines. Inlining is a well-known program transformation which substitutes procedure invocations with their corresponding implementations to erase artificial procedural boundaries [9]. However, conventional inlining cannot be applied to recursive procedures or subroutines defined in external libraries where the source code is not available. Further, when excessively applied, it can cause code size explosion and curtail the compiler’s effectiveness in applying optimizations (e.g., automatic loop parallelization and register allocation) due to the increased code complexity resulted from inlining.

This paper investigates techniques that enhance the effectiveness of inlining to support more aggressive loop parallelization by optimizing compilers. In particular, after identifying limitations of conventional inlining while using the Polaris compiler [7] to parallelize a collection of Fortran77 applications from the PERFECT benchmark suit [6], we present a new inlining approach to overcome these limitations. Our results show that the new approach can eliminate negative impact of conventional inlining while enhancing the effectiveness of interprocedural parallelization for a majority of the PERFECT benchmarks.

This research was supported in part by the National Science Foundation under Grants 0833203, 0747357, 0702527, 0855247, and 1027521 and by the Department of Energy under Grant DE-SC001770




Fig. 1. Workflow of our annotation-based inlining approach

Figure 1 shows the workflow of our enhanced inlining approach. In contrast to conventional inlining, which substitutes a procedure invocation with the complete implementation of the callee, we use annotations, which summarize the computational structure and side effects of the callee, to replace the invocation. The inlined code is then optimized by the Polaris compiler which applies sophisticated loop dependence analysis to automatically parallelize loops via OpenMP when safe. The optimized code from Polaris is then piped into a *reverse inliner*, which reverts the earlier inlined code back to using the original procedure invocations but leaves the OpenMP pragma intact. The output from the reverse inliner is essentially the original input code optimized with automatic parallelization, where the annotations, currently manually provided by developers, have been used to enable more aggressive loop optimization by the Polaris compiler in spite of opaque procedure calls. Our annotation-based inlining approach offers the following advantages over conventional inlining.

- Inlining can be applied even for subroutines defined in external libraries without their source code and for recursive subroutines because developers can provide a high-level summary of the semantics of these subroutines.
- The potential code size explosion problem can be avoided entirely as the inlining transformations will be reverted back to the original call statements after optimization.
- The user-supplied annotations do not need to include irrelevant implementation details of the subroutines of interest. After inlining, the code within the caller is much easier to analyze compared to when it contains the complete implementation of the callee.

The rest of this paper is organized as follows. Section II summarizes limitations of conventional inlining when used to support automatic loop parallelization by the Polaris compiler. Section III presents details of our annotation-based inlining approach. Section IV presents experimental results. Sections V and VI present related work and conclusions.

II. AUTOMATIC PARALLELIZATION USING POLARIS

Polaris is a source-to-source Fortran 77 compiler which supports automatic parallelization of loops based on sophisticated dependence analysis techniques [7]. It uses simple heuristics controlled via command-line options to govern whether each procedure call is inlined before parallelization analysis [17]. The default strategy inlines a procedure call only when the procedure contains no I/O and not many statements (≤ 150 by default) and when the invocation is inside a loop nest.

Polaris includes a number of sophisticated dependence analysis techniques which are fairly effective when analyzing regular Fortran DO loops operating on array subscripts that are linear combinations of the surrounding loop index variables. However, it becomes overly conservative when encountering non-linear array subscripts, which could be introduced by the inlining transformation applied before the analysis. The following summarizes the main issues we found to significantly hinder the effectiveness of Polaris loop parallelization analysis when combined with conventional procedure inlining.

A. Loss of Parallelism Due To Inlining

In languages such as Fortran/C/C++, arrays are treated as pointers into regions of data, and the same data operated by different subroutines can be declared as arrays of different shapes. Further, when optimizing Fortran subroutines, compilers can assume different array parameters are not aliased to each other. When subroutine invocations are inlined, the abstraction layer is broken, and the inlined implementations may become harder to analyze due to excessive code complexity introduced by inlining. As a result, loops that can be automatically parallelized by compilers when inside their respective subroutines may become no longer parallelizable by compilers after inlining, as discussed in the following.

1) *Forward Substitution of Non-linear Subscripts*: Figures 2-3 illustrate a situation where non-linear array subscripts are introduced by inlining the invocation of subroutine *PCINIT* at line 3 of Figure 3 with its implementation in Figure 2. Here the actual parameters used in the invocation are indirect references pointing to different regions of a global array *T*. When using these indirect array references to instantiate the formal parameters *X2*, *Y2*, and *Z2* of *PCINIT* in Figure 2, the array references *X2(I)*, *Y2(I)*, and *Z2(I)* at lines 8-10 of Figure 2 become *T(IX(7) + I)*, *T(IX(8) + I)* and *T(IX(9) + I)* respectively. Because the values of *IX(7)*, *IX(8)*, and *IX(9)* are unknown at compile time, the inlining transformation has created subscripted subscripts (array subscripts that contain additional subscripted array references) which are non-linear and considered non-analyzable by most dependence analysis techniques. As a result, the loops at lines 3 and 6 of Figure 2 can no longer be automatically parallelized after inlining, although Polaris dependence analysis can safely parallelize them inside the *PCINIT* subroutine before inlining.

2) *Linearization of Array Dimensions*: Figures 4-5 illustrate a situation where common arrays operated by two different subroutines are declared with different shapes. In particular, multi-dimensional arrays *PP*, *PHIT*, and *TM1* are used at line 5 of Figure 5 to invoke the subroutine *MATMLT* defined

```

1  SUBROUTINE PCINIT(X2,Y2,Z2,..)
2  DIMENSION X2(*),Y2(*),Z2(*)
3  .....
4  DO 200 N=1,NTYPES
5  NSP=NSPECI(N)
6  NS=NSITES(N)
7  DO 200 J=1,NSP
8  I=I+1
9  X2(I)=FX(I)*TSTEP**2/2.DO/DSUMM(N)
10 Y2(I)=FY(I)*TSTEP**2/2.DO/DSUMM(N)
11 Z2(I)=FZ(I)*TSTEP**2/2.DO/DSUMM(N)
12 .....
13 200 CONTINUE
14  ...

```

Fig. 2. A subroutine with automatically parallelizable loops at lines 3 and 6

```

1  DIMENSION T(*),S(*),W(*)
2  COMMON/WINDEK/ IX(99)
3  .....
4  CALL PCINIT(T(IX(7)),T(IX(8)),T(IX(9)),...)
5  .....

```

Fig. 3. A call site of the subroutine in Figure 2 (loops in *PCINIT* become no longer automatically parallelizable after inlining)

in Figure 4. However, the corresponding formal parameters *M1*, *M2*, and *M3* are declared as single-dimensional arrays in Figure 4. To inline the subroutine invocation, Polaris reconciles the mismatched array declarations by linearizing *PP*, *PHIT*, and *TM1* in Figure 5 into single dimensional arrays without any explicit shape information. After inlining, the compiler can no longer precisely recognize the dependence constraints of the inlined loops. As a result the three loops at lines 22, 23, and 26 of Figure 4 can no longer be parallelized after inlining.

B. Missed Opportunities

Conventional inlining substitutes a subroutine invocation with the entire implementation of the callee, where excessive complexity in the callee's implementation can force compiler optimizations, e.g., automatic loop parallelization, to be overly conservative due to the lack of domain-specific knowledge and

```

20  SUBROUTINE MATMLT(M1,M2,M3,L,M,N)
21  DOUBLE PRECISION M1(L*M),M2(M*N),M3(L*N)
22  DO 120 JL=1,L
23  DO 110 JN=1,N
24  J=JL+L*(JN-1)
25  M3(J)=0.0
26  DO 100 JM=1,M
27  K1=JL+L*(JM-1)
28  K2=JM+M*(JN-1)
29 100 M3(J)=M3(J)+M1(K1)*M2(K2)
30 110 CONTINUE
31 120 CONTINUE
32  RETURN
33  END

```

Fig. 4. A subroutine with automatically parallelizable loops at lines 22, 23, and 26

```

...
1  DIMENSION PP(4,4,15),PHIT(4,4),TM1(4,4),...
2  ...
3  DO 160 KS=1,15
4  KSM=KS-1
5  IF(KS.GT.1) CALL MATMLT(PP(1,1,KSM),PHIT,TM1,4,4,4)
6  ENDIF
7  ...
8 11 160 CONTINUE
9  ...

```

Fig. 5. A call site which invokes the subroutine in Figure 4 (loops in *MATMLT* are no longer automatically parallelizable after inlining)

```

1  SUBROUTINE FSMP(ID, IDE)
  ...
2  CALL GETCR(ID)
3  IRECT = IEGEOM(ID)
4  ICT   = IECURV(ID)
5  K1    = AK1(ICT)
6  K2    = AK2(ICT)
7  K12   = AK12(ICT)
8  ISTRES = 0
9  CALL SHAPE1
10 IF (IDEDON(IDE).EQ.0) THEN
11   IDEDON(IDE) = 1
12   CALL FORMF(FE(1,IDE))
13   CALL CHOFAC(FE(1,IDE), NSFE, IERR)
14   IF (IERR.NE.0) THEN
15     WRITE(6,*) ' F ELEMENT ',IDE,' IS SINGULAR '
16     STOP ' F SINGULAR '
17   ENDIF
18   CALL FORMS(SE(1,IDE))
19   CALL FORMM(ME(1,IDE))
20   CALL FORMNL(MNLE(1,IDE))
21 ENDIF
22 CALL GETLD(ID)
23 CALL FORMP(PE(1,ID))
24 RETURN
25 END

```

Fig. 6. A subroutine excluded from inlining by Polaris

```

1C . LOOP OVER THE SUBSTRUCTURES .
2  DO 35 ISS = 1, NSS
3C . LOOP OVER THE ELEMENTS IN THIS SUBSTRUCTURE .
4   DO 30 K = 1, NEPSS(ISS)
5C . FORM THE ELEMENTAL ARRAYS .
6    ID = IDBEGS(ISS) - 1 + K
7    IDE = K
8    CALL FSMP(ID, IDE)
9 30 CONTINUE
10 35 CONTINUE

```

Fig. 7. A loop nest invoking the subroutine in Figure 6 (the inner K loop at line 4 can be automatically parallelized after annotation-based inlining)

runtime information. The following discusses situations where the complexity of subroutine implementations prevents them from being effectively inlined or optimized.

1) *Opaque Compositional Subroutines*: Conventional inlining typically leaves out subroutines that make additional non-trivial procedure calls, as inlining a chain of subroutine invocations could result in serious code explosion. For example, the subroutine *FSMP* in Figure 6 is excluded from inlining by Polaris as it invokes a fair number of other subroutines. This subroutine serves to initialize a single column of five arrays, *FE* (lines 12-13), *SE* (line 18), *ME* (line 19), *MNLE* (line 20), and *PE* (line 23), using a large number of global variables, including both scalar variables and arrays, some of which are modified to hold intermediate results of the internal computation. In spite of the complexity of computation, distinct columns of the five arrays are modified when invoking *FSMP* with different values for *ID* and *IDE*. Figure 7 shows an example loop nest which invokes *FSMP* with values for *ID* obtained from a global array *IDBEGS* which returns a unique integer for each given value of *ISS*. After feeding such information to the Polaris compiler via annotations, the compiler is able to tell that distinct values of *ID* and *IDE* are used at different iterations of the inner K loop at line 4. As a result, it can automatically parallelize this loop after annotation-based inlining is applied, discussed in Section III-B2.

2) *Debugging and Error Checking*: In practical applications, debugging and error checking statements are often used inside subroutines to ensure proper termination of the appli-

```

1  SUBROUTINE GETCR(K)
  ...
2  COMMON /ECOORD/ XY(2,ZNNPED)
  ...
5 100 CONTINUE
6   DO 200 I = 1, NNPED
7     J = ICOND(I,K)
8     XY(1,I) = XYG(1,J)
9     XY(2,I) = XYG(2,J)
10 200 CONTINUE
  ...
11 RETURN
24 END

```

Fig. 8. Definition of subroutine *GETCR* invoked by *FSMP* in Figure 6

```

1  SUBROUTINE SHAPE1
  ...
2  COMMON /ECOORD/ XY(2,ZNNPED)
  ...
3  IF (NNPED.EQ.9) THEN
4   AX = XY(1,3) - XY(1,1)
5   AY = XY(2,5) - XY(2,3)
6  ELSE IF (NNPED.EQ.4) THEN
7   AX = XY(1,2) - XY(1,1)
8   AY = XY(2,3) - XY(2,2)
9  ELSE IF (NNPED.EQ.16) THEN
10  AX = XY(1,4) - XY(1,1)
11  AY = XY(2,7) - XY(2,4)
12 ELSE
13  CALL ERRORS(' NNPED NOT IMPLEMENTED IN SHAPE1')
14 ENDIF
  ...

```

Fig. 9. Definition of subroutine *SHAPE1* invoked by *FSMP* in Figure 6

cation when processing erroneous input data. This situation is illustrated by lines 14-17 of Figure 6, where the whole program would abort if previous evaluation has resulted in logical errors (indicated by the global *IERR* variable). Since debugging and error checking conditionals typically contain program I/O and early termination of the program, conservative compilers need to disable optimizations of the surrounding loops. However, pre-tested input data are often known to not trigger erroneous conditions at runtime, where the error handling statements are never executed. Even when errors do occur, replication of error messages is often acceptable. Such application-specific knowledge can be incorporated in our annotation-based inlining mechanism to support more aggressive loop parallelization, discussed in Section III-B3.

3) *Use of Temporary Arrays*: Many subroutines use temporary arrays to store intermediate results of computation, where each temporary array is first modified with new values before being used for additional computation. When the whole computation is inside a surrounding loop, compilers can apply array kill analysis to determine whether any value of the array comes from the previous iterations. If the whole array is killed (reinitialized) at each iteration, the temporary array can be privatized (duplicated within each thread) when parallelizing the surrounding loop. However, the array kill analysis may fail when only a subset of the array elements are modified, and those being used later are not obviously covered by the modifications. To illustrate such situations, Figures 8 and 9 provide two subroutine definitions invoked by the *FSMP* routine in Figure 6. Here a global array *XY* is used as a temporary array which is modified by the *GETCR* subroutine in Figure 8 and then used by the *SHAPE1* subroutine in Figure 9. Although *GETCR* modifies only a subset of the elements in *XY* (specifically, it modifies $XY(1:2,1:NNPED)$, where $NNPED \leq ZNNPED$), only those elements being

```

1  SUBROUTINE ASSEMR(ID, RHSE, RHSI, RHSB)
2  ...
3  DO 40 IN = 1, NNPD
4  NODE = ABS(ICOND(IN, ID))
5  IBLOCK = IWHERD(NODE, 1)
6  IREL = IWHERD(NODE, 2)
7  IF (IBLOCK.EQ.NBLOCK) THEN
8  DO 10 I = 1, NDDF
9  RHSB(IREL+I-1) = RHSB(IREL+I-1) + RHSE(I, IN)
10 CONTINUE
11 ELSE
12 DO 20 I = 1, NDDF
13 RHSI(IREL+I-1) = RHSI(IREL+I-1) + RHSE(I, IN)
14 CONTINUE
15 ENDIF
16 40 CONTINUE
17 RETURN

```

Fig. 10. A subroutine which contains indirect references in array subscripts

```

1  DO 10 K = 1, NEPSS(ISS)
2  ID = IDBEGS(ISS) + K - 1
3  IDE = IESMNO(ID)
4  CALL GETEU(ID, XE, X)
5  CALL MATMUL(ME(1, IDE), XE, MXE, NDFE, NDFE, 1)
6  CALL ASSEMR(ID, MXE, MXI, MXB)
7 10 CONTINUE

```

Fig. 11. A parallelizable loop invoking the subroutine in Figure 10

modified by *GETCR* are used in *SHAPE1*. However, due to the complexity of the multiple conditionals in Figure 9, a typical optimizing compiler would fail to discover the coverage relation even after both subroutines are successfully inlined. We resolve this issue by declaring that the whole temporary array is reinitialized via user-supplied annotations, shown in Figure 13 and discussed in Section III-B4.

4) *Indirect References In Array Subscripts*: Due to the lack of knowledge about runtime values of different arrays, conventional loop dependence analysis techniques are overly conservative when array references are used inside the subscripts of accessing other arrays. Figure 10 illustrates such an example, where two global arrays, *ICOND* and *IWHERD*, which serve to save one-to-one relations between data in different arrays, are used to compute the subscripts of modifying two other arrays *RHSB* and *RHSI*. After forward substitution of variables, the subscripts used to modify *RHSB* and *RHSI* at lines 6-13 become $IWHERD(ABS(ICOND(IN, ID), 2) + I - 1)$, which will always yield different values when given distinct values of *IN*, *ID* and *I*. Therefore different elements of the arrays *RHSB* and *RHSI* are modified when given distinct values of *IN*, *ID*, and *I*. However, such application-specific information is not available to the compiler, which must assume both *ICOND* and *IWHERD* could have arbitrary unknown values. Consequently the compiler must conservatively assume that arbitrary elements of *RHSB* and *RHSI* could be modified and must refrain from optimizing any loop that invokes *ASSEMR* even if the implementation of *ASSEMR* has been inlined. Figure 10 illustrates such a call site, where the *K* loop invokes *ASSEMR* at line 6 with the values of *ID* uniquely determined by the loop index variable *K*. We discuss how to enable Polaris to safely parallelize this loop via annotation-based inlining in Section III-B5.

III. ENHANCING THE ROLE OF INLINING

To improve the effectiveness of automatic loop parallelization when encountering situations discussed in Section II,

we seek to enhance the role of inlining so that higher-level semantics of subroutine invocations can be made readily available to compilers. In particular, we use the following steps to enhance the effectiveness of conventional inlining in supporting automatic parallelization across procedural boundaries.

- 1) Annotate important subroutines to summarize their side effects and loop structures required for accurate dependence analysis. Then, substitute subroutine invocations with the corresponding annotations instead of the actual detailed implementations.
- 2) Use Polaris to apply conventional loop dependence analysis and automatically parallelize loops when safe by inserting OpenMP directives.
- 3) Apply a *reverse inlining* step which substitutes all the inlined annotations with appropriate invocations of the original subroutines. After this step, the only remaining transformation to the original input code is the parallelization of loops via OpenMP.

The reverse inlining step essentially reverses all the transformations introduced by annotation-based inlining so that the original input code can benefit from advanced compiler optimizations without sacrificing its modularity. The correct application of this step requires all the inlined annotations be recognized and mapped back to correct invocations of the original subroutines, which can be easily accomplished when minimal transformations, e.g., insertion of OpenMP directives, have been made to the inlined code. However, the task is more challenging when interacting with more drastic program transformations, e.g., loop blocking and software pipelining. Section III-C discusses these issues in more detail.

The following first discusses our annotation language and then illustrates how to use annotations to summarize the higher-level semantics of subroutines and enable more effective parallelization after inlining. Section III-C presents our algorithm for enhanced inlining. Section III-D discusses the correctness and generality of the overall approach.

A. The Annotation Language

Figure 12 summarizes the syntax of our annotation language, which can be used by developers to describe the side effects and loop structures of important subroutines. When these annotations are used to support subroutine inlining, a compiler can correctly recognize the dependence constraints carried by each subroutine invocation and subsequently successfully parallelize surrounding loops when safe.

The statements supported by our language include loops, if-conditionals, assignments, variable declarations, and return statements. They are used to summarize the control-flow structure and memory side effects of each subroutine. For implementation details that cannot be expressed using these statements, we provide two special operators, *unique* and *unknown*, to summarize approximate relations between variables while omitting details of the computation. In particular, $y = \text{unique}(x_1, \dots, x_n)$ specifies that the value of *y* is uniquely computed (determined) from those of variables x_1, \dots, x_n ; that is, if y_1 is computed from $(x_1=v_1, \dots, x_n=v_n)$, y'_1 is computed from $(x_1=v'_1, \dots, x_n=v'_n)$, and $(v_1, \dots, v_n) \neq (v'_1, \dots, v'_n)$, then

```

s : {s1 s2 ... sn}
  | if (e) s1 [else s2]
  | do (id=e1:e2[:e3]) s
  | var=e;
  | vars=unknown(e1, e2, ...);
  | vars=unique(e1, e2, ...);
  | type var1,...,varn;
  | return e;
var : id | id '[' e1, ..., en ']'
vars : var | (var, ..., var)

s, s1, ..., sn : statements;
e, e1, ..., en : expressions;
id : variable name;
id [ e1, ..., en ] : multi-dimensional array reference;
type var1,...,varn : declaring types of variables;

```

Fig. 12. The annotation language

$y_1 \neq y'_1$. Therefore, if the values of x_1, \dots, x_n are different at distinct iterations of a loop surrounding $y=unique(x_1, \dots, x_n)$, then the values of y are guaranteed to be similarly different. In contrast, $y=unknown(x_1, \dots, x_n)$ specifies that the value of variable y is computed from reading those of variables x_1, \dots, x_n , but the relationship could be arbitrary. These special-purpose operators serve to abstract away complex implementation details which degrade the effectiveness of compiler analysis, while keeping essential relations among variables visible to the compiler.

Expressions supported by our annotation language include most of the arithmetic operations in Fortran 77 combined with memory references via scalar and array variables. The Fortran 90 notation of array regions are supported so that collective operations can be applied to arrays without requiring explicit loops. The two special operators, *unknown* and *unique*, can also be used directly inside expressions, where their results do not need to be saved in variables before used.

B. Writing Annotations

Our annotation language serves to accurately summarize the side effects and loop structures of subroutines without exposing their local implementation details that are irrelevant to the surrounding calling context. In particular, for each subroutine, the annotations aim to summarize relations between its input parameters and output values as well as global variables modified by the subroutine while omitting intermediate results and variables that are local to the subroutine. The goal is to minimize adverse side effects of conventional inlining which may result in accidental loss of parallelism in the inlined code. To demonstrate the capacity of this approach, the following illustrates how to use user-supplied annotations to overcome inefficiencies of automatic parallelization discussed in Section II.

1) *Avoiding Loss Of Parallelism*: As discussed in Section II-A, when conventional inlining breaks procedural boundaries by substituting subroutine invocations with detailed implementations, some parallel loops inside the inlined code may become no longer parallelizable by compilers due to unexpected interactions between the caller and callee. Our annotation-based approach resolves this issue by preserving all the original procedural boundaries and thereby entirely eliminating the adverse side effects of conventional inlining. In particular, after we enable the compiler to perform interprocedural dependence analysis by substituting subroutine

```

subroutine FSMP (ID, IDE) {
  XY=unknown (XYG [:, ICOND [:, ID]], NSYMM);
  IRECT=IEGEOM [ID];
  K1=AK1 [IECURV [ID]];
  K2=AK2 [IECURV [ID]];
  K12=AK12 [IECURV [ID]];
  ISTRES=0;
  (NDX, NDY, WTDET)=unknown (IRECT, XY, NDXI, NDETA, QDWGHT,
                             NQD, NNPED);
  if (IDEDON [IDE]==0) {
    IDEDON [IDE]=1;
    FE [:, IDE]=unknown (WTDET, C, NB, NNPE, NQD, NSF);
    ME [:, IDE]=unknown (WTDET, N, NSYMM, RHO, NQD, NNPED);
    SE [:, IDE]=unknown (WTDET, NB, N, NDX, NDY, K1, K2, K12,
                        NNPE, NQD, NNPED);
    MNLE [:, IDE]=unknown (WTDET, NB, NDX, NDY, KNONLN,
                          NNPE, NQD, NNPED);
  }
  P=unknown (PXYZ [:, ABS (ICOND [:, ID])], NNPED, GX, NSYMM);
  PE [:, ID]=unknown (P, WTDET, N, NQD, NNPED);
}

```

Fig. 13. Annotations for the opaque subroutine defined in Figure 6

invocations with user-supplied annotations, the reverse inlining step will replace the inlined annotations with the original subroutine invocations, thereby preserving the original procedural boundaries as well as their optimizations.

2) *Summarizing Opaque Subroutines*: Since we substitute subroutine invocations with summaries of their semantics supplied by developers, our approach can be easily applied to opaque subroutines with arbitrarily complex implementations. As example, Figure 13 illustrates our annotations for the *FSMP* subroutine in Figure 6, which was excluded from inlining by Polaris due to its excess code complexity. In particular, these annotations summarize the regions of global arrays (*FE*, *ME*, *SE*, *MNLE*, *PE*) modified by all the subroutines invoked from *FSMP*, the temporary global variables (*XY*, *IRECT*, *K1*, *K2*, *K12*, *ISTRES*, *NDX*, *NDY*, *WTDET*) modified in the process, and the read-only global variables (*IEGEOMI*, *IECURV*, among others) used in the computation. The *unknown* operator is used extensively to omit local implementation details (e.g., intermediate results) of relevant computation, allowing invocations of the subroutine to be more accurately handled by dependence analysis of their surrounding loops.

3) *Debugging and Error Handling*: Many subroutines in large applications contain program output statements used for debugging and error handling purposes. The presence of these *exception handling* statements are typically treated with extreme caution by compilers, where all reordering transformations of the surrounding loops are subsequently disabled. However, since these statements are used for debugging/error handling only, they are not triggered at runtime in most cases, and even when triggered, precise exception handling is often not required. Using our annotation-based inlining approach, developers can choose to relax the consistency requirement of exception handling when parallelizing their program, by omitting these situations in the subroutine annotations. For example, in Figure 13, the error checking conditional at lines 14-17 of Figure 6 has been omitted in the annotations. Therefore it no longer prevents loops surrounding invocations of the *FSMP* subroutine from being safely parallelized.

4) *Use Of Temporary Arrays*: Complex subroutines often use temporary arrays to hold intermediate results of the internal computation. When these arrays are declared as local variables, our annotations will omit their existence entirely

```

subroutine ASSEMR(ID,RHSE, RHSI, RHSB) {
  do (IN=1:NNPED)
    do (I=1:NDDF)
      if (unique(ID,IN) == NBLOCK)
        RHSB[unique(ID,IN,I)] += RHSE[I,IN];
      else
        RHSI[unique(ID,IN,I)] += RHSE[I,IN];
    }
}

```

Fig. 14. Annotations for the subroutine in Figure 10

as they do not incur any visible side effects to the outside. However, sometimes these arrays are declared in the global scope and used to pass values from one subroutine to another. An example global array used for this purpose is shown in Figure 8, where the global array *XY* is modified in subroutine *GETCR* to hold intermediate results and then used in the subroutine *SHAPE1* in Figure 9. It is conceptually a temporary array within the *FSMP* subroutine as only those elements defined in *GETCR* are used in the subsequent calls to *SHAPE1* and other subroutines. Similar global temporary arrays in *FSMP* include *NDX*, *NDY*, *WDET* and *P*, shown in Figure 13. In our annotations, these arrays are modified and used as if they are atomic scalar variables. Since modifications to these variables precede all their uses in the annotations of subroutine *FSMP* in Figure 13, they can be treated as private variables when parallelizing a loop surrounding the invocation of *FSMP*, shown in Figure 7. In particular, when parallelizing the *K* loop in Figure 7, Polaris would peel the last iteration of the loop before parallelizing all the other iterations by privatizing temporary arrays in those iterations, so all the global arrays have the same values as their original sequential computation after the entire loop is finished.

5) *Indirect References In Array Subscripts*: Many Fortran applications use global arrays to store dynamic relations between different data structures. Most of these arrays are initialized only once throughout the entire program to save a one-to-one unique mapping between the related data. An example subroutine using these special-purpose arrays is shown in Figure 10, where *ICOND* and *IWHERD* are global arrays which contain one-to-one relations between elements stored in different places. When they are used as subscripts to access *RHSB* and *RHSI* in Figure 10, the subscript expressions are non-linear, and compilers have to be overly conservative when parallelizing the surrounding loops. To overcome these difficulties, in Figures 14, we use the *unique* operator to summarize the values of these arrays in terms of the relevant input parameters and loop index variables (*ID*, *IN*, and *I*). The declaration of the *unique* relation comes from domain-specific knowledge of the developer. When using the annotations in Figure 14 to substitute for the invocation of subroutine *ASSEMR* in Figure 11, each unique operator will be replaced with a linear expression which uniquely combines the involved integer variables *ID*, *IN*, and *I*. As a result the compiler can easily recognize that unique elements of arrays *RHSB* and *RHSI* are modified at each distinct iteration of the surrounding loop and thereby can safely parallelize the loop in Figure 11.

C. The Enhanced Inlining Algorithm

Figure 15 shows the key steps of our algorithm for applying automatic parallelization with enhanced inlining support. The

Input: program source code and annotations for selected subroutines

Output: optimized source code

Algorithm:

- 1) Annotation-based inlining: For each call statement within the input program where annotations are provided for the callee
 - a) Instantiate the corresponding annotations with actual parameters;
 - b) Replace the call statement with the instantiated annotations;
 - c) insert tags surrounding the instantiated code fragment from inlining.
- 2) Automatic parallelization: invoke Polaris to optimize the inlined code.
- 3) Reverse inlining: For each *tagged* code segment in the optimized source code:
 - a) Find the corresponding subroutine annotations from looking into the tag;
 - b) Match the tagged code segment against the corresponding annotations to compute instantiation parameters;
 - c) Replace the tagged code segment with a subroutine call using the instantiation parameters.

Fig. 15. Automatic parallelization with annotation-based inlining

algorithm is comprised of three main phases: annotation-based inlining, automatic parallelization, and reverse inlining. The following explains each phase in detail.

1) *Annotation-based Inlining*: The implementation of this step is similar to conventional inlining, except that subroutine invocations are substituted with user-supplied annotations instead of detailed implementations of the callees. Translating annotations to the underlying programming language (e.g., Fortran) is trivial, except for the two special purpose operators, *unknown* and *unique*. To translate each *unknown* operator, we define a new uninitialized global array, modify the array with all the operands of the *unknown* operator, and then replace the *unknown* invocation with an access to the new array. To translate each *unique* operator, we replace it with a linear expression which uniquely combines all the relevant integer variables. After inlining, a pair of special tags are placed surrounding the inlined source code to support reverse inlining at a later stage. Figure 18 shows the result of applying annotation-based inlining to an invocation of the *MATMLT* subroutine in Figure 5, using annotations in Figure 16.

In our algorithm, only subroutines with annotations are considered for inlining. Note that although the original implementation of *MATMLT* in Figure 4 declares the array parameters *M1*, *M2*, *M3* as having single dimensions, our annotations declare them as two-dimensional matrices. Subsequently our annotation-based inlining can avoid the unnecessary linearization of array dimensions which may degrade the precision of compiler analysis, as discussed in Section II-A.

2) *Automatic parallelization*: After applying annotation-based inlining to the input program source, we optimize the inlined code using the Polaris compiler (with conventional inlining disabled), which performs advanced dependence analysis of the inlined code and automatically inserts OpenMP directives to parallelize loops when safe and profitable (the profitability is determined based on simplistic heuristics, e.g., all parallelized loop needs to exceed a certain number of iterations). Figure 17 shows the result of applying loop parallelization to our inlined code in Figure 18. Note that the pair of special tags surrounding the inlined annotations remain intact after the parallelization optimizations.

3) *Reverse-inlining*: After automatic parallelization, the reverse-inlining step is performed to reverse the annotation-based inlining transformation while keeping the OpenMP directives inserted by the Polaris compiler intact. This step is necessary to ensure correctness of program optimization.

```

subroutine MATMLT(M1,M2,M3,L,M,N) {
  dimension M1[L,M], M2[M,N], M3[L,N];
  M3 = 0.0;
  do (JN=1:N)
    do (JM=1:M)
      M3[:,JN] += M1[:,JM] * M2[JM,JN];
    }
}

```

Fig. 16. Annotations for the *MATMLT* subroutine in Figure 4

```

...
!$OMP PARALLEL
!$OMP+DEFAULT(SHARED)
!$OMP DO
  DO KS=1,15
    IF(KS.GT.1) THEN
* //@; BEGIN(Code)
* @annot inline MATMLT {
!$OMP PARALLEL
!$OMP+DEFAULT(SHARED)
!$OMP DO
  DO JL=1,4,1
  DO JN=1,4,1
    TM1(JL,JN)=0.0
    DO JM=1,4,1
      TM1(JL,JN)=TM1(JL,JN)
*+PP(JL,JM,KS-1)*PHIT(JM,JN)
    ENDDO
  ENDDO
  ENDDO
!$OMP END DO NOWAIT
!$OMP END PARALLEL
* @}
  ENDDO
...
!$OMP END DO NOWAIT
!$OMP END PARALLEL

```

Fig. 17. Call site of *MATMLT* in Figure 18 after parallelization

Specifically, while the user-supplied annotations are expected to carry equivalent side effects and dependence constraints as the real subroutine implementation, they are typically not semantically equivalent to the original implementation, due to simplification of internal implementation details and the use of the special-purpose summary operations, e.g., the *unknown* and *unique* operators. Therefore, the inlined annotations must be reversed back to an equivalent subroutine invocation to guarantee the correctness of the optimized code.

The reverse inlining transformation is applied to all the tagged code segments created by the earlier annotation-based inlining transformation. For each tagged fragment, it first finds the corresponding subroutine annotations and then proceeds to compute an instantiation value for each formal parameter of the subroutine. Specifically, when using these parameter values to instantiate the subroutine annotations, the resulting code must be equivalent to the tagged code segment. Currently we apply a pattern matching algorithm to compare the tagged code segment and the inlining annotations node by node, while allowing variable substitution, expression reordering, and OpenMP directives inside the tagged segment. Since the only optimization performed by Polaris is the insertion of OpenMP directives, which are simply ignored in the pattern matching process, a set of appropriate parameter values are guaranteed to be found for each tagged code segment. These values are then used as actual parameters to generate a subroutine invocation so that each tagged code segment is replaced with its original function call.

Figure 19 shows the resulting code after applying reverse

```

...
DIMENSION PP(4,4,15),PHIT(4,4),TM1(4,4),...
...
DO KS=1,15
  KSM=KS-1
  IF(KS.GT.1) THEN
*//@; BEGIN(Code)
  @annot inline MATMLT {
    DO JL=1,4,1
      DO JN=1,4,1
        TM1(JL,JN)=0.0
        DO JM=1,4,1
          TM1(JL,JN)=TM1(JL,JN)
*PP(JL,JM,KSM)*PHIT(JM,JN)
        ENDDO
      ENDDO
    ENDDO
  @}
  ENDDO
...
ENDDO
...

```

Fig. 18. Call site of *MATMLT* in Figure 5 after annotation-based inlining

```

...
!$OMP PARALLEL
!$OMP+DEFAULT(SHARED)
!$OMP DO
  DO KS=1,15,1
    IF(KS.GT.1) THEN
      CALL MATMLT(PP(1,1,KS-1),PHIT(1,1),TM1(1,1),4,4,4)
    ENDDO
  ENDDO
!$OMP END DO NOWAIT
!$OMP END PARALLEL
...

```

Fig. 19. Call site of *MATMLT* in Figure 17 after reverse-inlining

inlining to the parallelized loop in Figure 17. Note that while the single optimization applied by Polaris is the insertion of OpenMP directives, the compiler does perform several normalization transformations, e.g., reordering of statements, induction variable substitution, and constant propagation, to the tagged code segments. As a result our reverse inlining transformation cannot simply replace them with the original subroutine calls. Our pattern matching algorithm is tolerant of minor modifications to the inlined annotations and can automatically extract the correct actual parameters in subroutine invocation in spite of the normalization transformations.

D. Correctness, Efficiency, and Generality

The correctness of our enhanced inlining approach depends on the soundness of the user-supplied annotations. Specifically, if the annotations accurately summarize the side-effects and dependence constraints of the subroutines, the automatic parallelization optimization is guaranteed to be safe. Currently, such consistency is not automatically verified, and we use runtime testers to check and verify the correctness of our optimized code. Our future work will develop techniques to automatically verify the soundness of user-supplied annotations and to automatically generate inlining annotations when possible.

The compile-time overhead of applying annotation-based inlining is similar to that of conventional inlining except that since user-supplied annotations are expected to be much shorter than detailed implementations, the cost of applying annotation-based inlining is lower, and the inlined annotations are expected to be much easier to analyze by the compiler.

| Applications | Descriptions |
|--------------|---|
| ADM | Pseudospectral air pollution simulation |
| ARC2D | Two-dimensional fluid solver of Euler equations |
| FLO52Q | Transonic inviscid flow past an airfoil |
| OCEAN | Two dimensional ocean simulation |
| BDNA | Molecular dynamic package for the simulation of nucleic acids |
| MDG | Molecular dynamics for the simulation of liquid water |
| QCD | Quantum chromodynamics |
| TRFD | A kernel simulating a two-electron integral transformation |
| DYFESM | Structural dynamics benchmark (finite element) |
| MG3D | Depth migration code |
| TRACK | Missile tracking |

TABLE I
SUMMARY OF THE PERFECT BENCHMARKS

Since all the inlined annotations are reversed back to using the original subroutine invocations, the modularity of the original program is not affected. The cost of applying reverse inlining to each tagged code segment is proportional to the size of the corresponding subroutine annotations, with constant cost associated with tolerating local modifications of the inlined annotations.

Our enhanced inlining approach can potentially enable general-purpose compilers to better utilize domain-specific knowledge from developers in supporting more effective inter-procedural optimization of large-scale applications. So far we have used this approach to support only automatic parallelization via OpenMP. When applying pattern-matching to reverse inlined code segments back to appropriate subroutine calls, our reverse inlining transformation can tolerate local modifications to the inlined code such as reordering of expressions, induction variable substitutions, and insertion of OpenMP directives. However, to extend our approach to similarly support other optimizations such as loop blocking and unrolling, it may become much more challenging to reverse the inlined code segments back to appropriate subroutine calls after dramatic modifications to the tagged segments. Therefore to apply this approach more extensively in a general-purpose compiler, a more systematic approach needs to be developed to compute correct instantiation parameters after dramatic modifications to the inlined annotations, which is a subject of our future work.

IV. EXPERIMENTAL RESULTS

To evaluate the effectiveness of our enhanced inlining approach when used to enable more aggressive automatic parallelization of loops, we selected 12 applications from the *PERFECT* benchmark suite [27], summarized in Table I. For each benchmark, we applied both conventional inlining and our enhanced inlining combined with automatic loop parallelization by the Polaris compiler. To measure the amount of parallelism enabled by and the degree of code explosion resulted from inlining, we counted the number of loops being parallelized after optimization and the line numbers of the resulting source code. We then use two multi-core machines, an Intel Macintosh running MacOS 10.5 with two quad-core 3GHz Intel processors (32KB L1 cache per core) and an AMD Opteron running Linux with two dual-core 3GHz AMD Opteron processors (128KB L1 cache per core), to measure the performance of the parallelized code. All benchmarks are compiled using *gfortran* 4.2.1 on the Intel Mac and *iFort* 11.1 on the AMD Opteron, using the `-O3` optimization flag.

A. Enhancing Automatic Loop Parallelization

Table II compares the number of automatically parallelized loops by the Polaris compiler using three different inlining configurations: disable inlining of all subroutines (i.e., *no inlining*); inline implementations of subroutines with ≤ 150 lines of source code via *conventional inlining* (the default inlining strategy adopted by Polaris); and *annotation-based inlining*, where our enhanced inlining approach is used. For each configuration, we counted the number of automatically parallelized loops (`#par-loops`) and the overall code size (the number of source code lines with all comments removed) after optimization by Polaris. Note that when conventional inlining is applied, Polaris could fail to parallelize some loops which were parallelizable when *no inlining* is applied, as discussed in Section II-A. These loops are categorized as `#par-loss` in Table II. Inlining may also enable additional loops being parallelized beyond those parallelized using *no-inlining*, these loops are categorized as `#par-extra` in Table II. If parallelized, each loop in the original benchmark is counted only once, even when inlining has made multiple copies of the original loop and all copies are subsequently parallelized.

From table II, inlining (including both conventional inlining and annotation-based inlining) is able to improve the effectiveness of automatic parallelization for 6 out of the 12 *PERFECT* benchmarks. For the other benchmarks, Polaris was not able to identify additional parallelism from loops which contain subroutine calls. Note that we have manually annotated a subset of subroutines from the *PERFECT* benchmarks based on careful inspection of their implementations. It may be possible to parallelize more loops by annotating additional subroutines. Our future work will investigate more systematic application of our enhanced inlining approach by automatically generating annotations. We have verified the correctness of all the automatically parallelized loops via both manual inspection and runtime testing of the parallelized code.

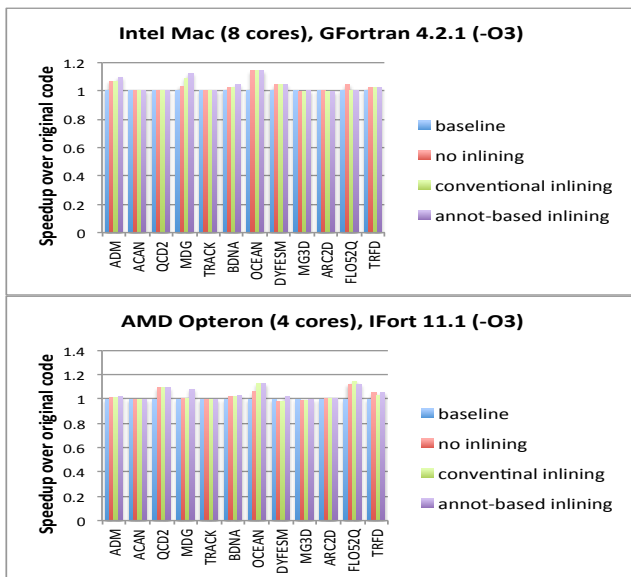
When combined with annotation-based inlining, the Polaris compiler is able to identify 37 additional parallelizable loops in different *PERFECT* benchmarks when compared with no-inlining. Most of these loops invoke complex subroutines which in turn invoke other routines, and the complexity of their implementations would overwhelm most state-of-the-art program analysis techniques. Since our annotation-based inlining allows developers to intervene with their application-specific knowledge, we are able to summarize the intended semantics of these subroutines to enable more effective parallelization of their surrounding loops. Examples of such annotations are illustrated in Section III-B. When applying annotation-based inlining, the code explosion problem is avoided entirely, as the reverse inlining step has restored all the original subroutine invocations (the small increase in code size is mostly due to the extra OpenMP directives inserted to parallelize loops).

In contrast, conventional inlining enabled Polaris to parallelize only a small subset (12 out of 37) of the extra parallel loops identified by annotation-based inlining. Additionally, after conventional inlining, Polaris can no longer parallelize 90 loops which were categorized as parallelizable when no inlining is performed, due to issues discussed in Section II-A.

| Applications | Total # of loops | No inlining | | Conventional inlining | | | | Annotation-based inlining | | | |
|--------------|------------------|-------------|-----------|-----------------------|------------|-----------|-----------|---------------------------|------------|-----------|-----------|
| | | # par-loops | code size | # par-loops | #par-extra | #par-loss | code size | # par-loops | #par-extra | #par-loss | code size |
| ADM | 268 | 179 | 8261 | 153 | 3 | 29 | 8703 | 190 | 11 | 0 | 8380 |
| ACAN | 147 | 30 | 36774 | 30 | 0 | 0 | 36774 | 30 | 0 | 0 | 36797 |
| QCD2 | 157 | 102 | 3498 | 102 | 0 | 0 | 4795 | 110 | 8 | 0 | 3574 |
| MDG | 52 | 37 | 1962 | 41 | 6 | 2 | 2459 | 43 | 6 | 0 | 2109 |
| TRACK | 87 | 54 | 3107 | 54 | 0 | 0 | 3465 | 55 | 1 | 0 | 3121 |
| BDNA | 219 | 129 | 6713 | 128 | 1 | 2 | 6371 | 134 | 5 | 0 | 6722 |
| OCEAN | 133 | 106 | 8607 | 106 | 0 | 0 | 8691 | 106 | 0 | 0 | 8631 |
| DYFESM | 197 | 133 | 4713 | 134 | 2 | 1 | 5674 | 139 | 6 | 0 | 4772 |
| MG3D | 150 | 51 | 22878 | 51 | 0 | 0 | 30100 | 51 | 0 | 0 | 22918 |
| ARC2D | 208 | 182 | 5185 | 139 | 0 | 43 | 4931 | 182 | 0 | 0 | 5290 |
| FLOS2Q | 175 | 149 | 3547 | 149 | 0 | 0 | 3562 | 149 | 0 | 0 | 3551 |
| TRFD | 38 | 27 | 1703 | 14 | 0 | 13 | 636 | 27 | 0 | 0 | 1710 |
| Totals | 1831 | 1179 | 106948 | 1101 | 12 | 90 | 116161 | 1216 | 37 | 0 | 107575 |

*Code size is computed as the number of source code lines with all comments removed.

TABLE II
AUTOMATICALLY PARALLELIZED LOOPS USING DIFFERENT INLINING STRATEGIES



* Baseline: performance of the original benchmarks with no optimization;

Fig. 20. Performance of automatically parallelized code

After conventional inlining, the code size increased by about 10% even when only small subroutines were inlined. The increase is likely significantly higher when more extensive inlining is applied via the conventional approach.

B. Performance of Optimized Code

Figure 20 presents the runtime speedups achieved by the automatically parallelized benchmarks when using different inlining configurations. Note that a majority of the PERFECT benchmarks do not benefit from loop parallelization due to their small input data size, which is a known issue for these benchmarks [25]. To avoid degradation of performance by excessive parallelization of loops, we used empirical performance tuning to disable a selected set of loops from being parallelized if their parallelization incurs a slowdown of the overall execution time. As demonstrated in Figure 20, at most 10% performance improvement is achieved by automatic loop parallelization combined with different inlining configurations. Annotation-based inlining is able to achieve the best performance for two benchmarks (ADM and MDG) and has

achieved similar performance as other inlining configurations for the other benchmarks.

V. RELATED WORK

Inlining is a widely adopted technique which can be used by compilers to erase procedural boundaries and apply optimizations to larger regions of code [1], [4], [30]. Ayers et. al [4] shown that aggressive inlining and cloning based on profiled information can dramatically improve the effectiveness of a large number of back-end optimizations. However, when excessively applied, inlining can cause the well-known code-explosion problem [13] and could degrade the effectiveness of many compiler optimizations as the input code becomes overwhelmingly large and complex. Previous research has developed a variety of heuristics, including temperature heuristics [36], demand-driven online transformation [32], inlining trials [14], and interprocedural flow analysis [3], [23], to selectively apply inlining so that performance benefits can be gained without incurring serious problems [2]. Other directions include automatic tuning of different inlining heuristics [10], combining inlining with hot code outlining [37], region-based compilation [21], [24], and whole program optimization by merging and re-dividing code regions [31].

This paper proposes an annotation-based inlining approach to overcome the negative impact of conventional inlining both in terms of code explosion and in terms of unexpected complications. Our use of developer-supplied annotations to summarize the semantics of opaque subroutines is similar to the *semantic inlining* approach by Wu *et al.*[35], [34], which allows their compiler to treat user-defined abstractions as primitive types in Java. The Broadway[18] and DyC[16] compilers used *annotation languages* to guide domain-specific optimizations and dynamic compilation of C code. Annotations and semantic specifications have also been used extensively to specify dynamic properties of lower-level implementations in program verification [15], [11]. We focus on using annotations to substitute for implementations of subroutines in inlining to enable more aggressive automatic parallelization.

Optimizing compilers have a long history of supporting automatic parallelization of user applications [33], [29], [26], [8], [20], [19], [12]. This paper focuses on using inlining to enable interprocedural parallelization without resorting to

expensive inter-procedural program analysis techniques [5], [22], [28]. Our work is orthogonal to and can be integrated with existing other compiler frameworks for automatic loop parallelization besides the Polaris compiler.

VI. CONCLUSIONS AND FUTURE WORK

This paper presents a study which exposes some serious limitations of conventional inlining when using the Polaris compiler [7] to parallelize a collection of Fortran applications from the PERFECT benchmark suite [6]. We then present a new annotation-based inlining approach to overcome the limitations. Our experimental results show that the new approach can eliminate most of the negative impact of conventional inlining while significantly enhancing the effectiveness of automatic loop parallelization across procedural boundaries. Our future work will develop techniques to automatically derive necessary annotations and to verify the safety of manually supplied annotations.

REFERENCES

- [1] R. Allen and K. Kennedy. *Optimizing Compilers for Modern Architectures*. Morgan Kaufmann, San Francisco, October 2001.
- [2] M. Arnold, S. Fink, V. Sarkar, and P. F. Sweeney. A comparative study of static and profile-based heuristics for inlining. *SIGPLAN Not.*, 35(7):52–64, 2000.
- [3] J. M. Ashley. The effectiveness of flow analysis for inlining. In *In Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming*, pages 99–111, 1997.
- [4] A. Ayers, R. Schooler, and R. Gottlieb. Aggressive inlining. *SIGPLAN Not.*, 32(5):134–145, 1997.
- [5] V. Balasundaram and K. Kennedy. A technique for summarizing data access and its use in parallelism enhancing transformations. *SIGPLAN Not.*, 24(7):41–53, 1989.
- [6] W. Blume and R. Eigenmann. Performance analysis of parallelizing compilers on the perfect benchmarks programs. *IEEE Transactions on Parallel and Distributed Systems*, 3:643–656, 1992.
- [7] W. Blume, R. Eigenmann, K. Faigin, J. Grout, J. Hoeflinger, D. Padua, P. Petersen, W. Pottenger, L. Rauchweger, P. Tu, and S. Weatherford. Polaris: Improving the effectiveness of parallelizing compilers. In *In Languages and Compilers for Parallel Computing*, pages 141–154. Springer-Verlag, 1994.
- [8] U. Bondhugula, A. Hartono, and J. Ramanujam. A practical automatic polyhedral parallelizer and locality optimizer. In *In PLDI 08: Proceedings of the ACM SIGPLAN 2008 conference on Programming language design and implementation*, 2008.
- [9] B. Calder and D. Grunwald. Reducing indirect function call overhead in c++ programs. In *POPL '94: Proceedings of the 21st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 397–408, New York, NY, USA, 1994. ACM.
- [10] J. Cavazos and M. F. P. OBoyle. Automatic tuning of inlining heuristics. In *IN ACM/IEEE CONFERENCE ON SUPERCOMPUTING*, page 14, 2005.
- [11] S. Chaki, E. Clarke, and A. Groce. Modular verification of software components in c. *Transactions of Software Engineering*, 1(8), Sept. 2004.
- [12] Y.-S. Chang, H.-J. Lee, D.-S. Park, and I.-Y. Lee. Interprocedural transformations for extracting maximum parallelism. In *ADVIS '02: Proceedings of the Second International Conference on Advances in Information Systems*, pages 415–424, London, UK, 2002. Springer-Verlag.
- [13] K. D. Cooper, M. W. Hall, and L. Torczon. Unexpected side effects of inline substitution: a case study. *ACM Lett. Program. Lang. Syst.*, 1(1):22–32, 1992.
- [14] J. Dean and C. Chambers. Towards better inlining decisions using inlining trials. In *In Proceedings of the 1994 ACM Conference on LISP and Functional Programming*, pages 273–282, 1994.
- [15] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation*, PLDI '02, pages 234–245, New York, NY, USA, 2002. ACM.
- [16] B. Grant, M. Mock, M. Philipose, C. Chambers, and S. J. Eggers. Dyc: an expressive annotation-directed dynamic compiler for c. *Theor. Comput. Sci.*, 248(1-2):147–199, 2000.
- [17] J. R. Grout. Inline expansion for the polaris research compiler. Technical report, 1995.
- [18] S. Z. Guyer and C. Lin. Broadway: A compiler for exploiting the domain-specific semantics of software libraries. *Proceedings of the IEEE, Special issue on program generation, Optimization, and Adaptation*, 93(2):342–357, 2005.
- [19] M. Hall, K. Kennedy, and K. S. M. Kinley. Interprocedural transformations for parallel code generation. In *In Proceedings of Supercomputing '91*, pages 424–434, 1991.
- [20] M. W. Hall, J. M. Mellor-Crummey, A. Carle, and R. G. Rodriguez. Fiat: A framework for interprocedural analysis and transformation. In *IN PROCEEDINGS OF THE SIXTH WORKSHOP ON LANGUAGES AND COMPILERS FOR PARALLEL COMPUTING*, pages 522–545. Springer-Verlag, 1995.
- [21] R. E. Hank, W. mei W. Hwu, and B. R. Rau. Region-based compilation: An introduction and motivation. In *in Proceedings of the 28th Annual International Symposium on Microarchitecture*, pages 158–168, 1995.
- [22] J. P. Hoeflinger, Y. Paek, and K. Yi. Unified interprocedural parallelism detection. *Int. J. Parallel Program.*, 29(2):185–215, 2001.
- [23] S. Jagannathan and A. Wright. Flow-directed inlining. In *In Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 193–205, 1996.
- [24] Y. Liu, Y. L. Zhaoqing, R. Qiao, and R. D. ching Ju. A region-based compilation infrastructure. In *In Proc. of the 7th Workshop on Interaction between Compilers and Computer Architectures*, pages 75–84. IEEE Computer Society Press, 2003.
- [25] K. Psarris and K. Kyriakopoulos. The impact of data dependence analysis on compilation and program parallelization. In *ICS '03: Proceedings of the 17th annual international conference on Supercomputing*, pages 205–214, New York, NY, USA, 2003. ACM.
- [26] S. Ramaswamy, S. Sapatnekar, and P. Banerjee. A framework for exploiting data and functional parallelism on distributed memory multicomputers, 1994.
- [27] R. H. Saavedra and A. J. Smith. Analysis of benchmark characteristics and benchmark performance prediction. *ACM Trans. Comput. Syst.*, 14:344–384, November 1996.
- [28] H. Seidl and B. Steffen. Constraint-based inter-procedural analysis of parallel programs. *Nordic J. of Computing*, 7(4):375–400, 2000.
- [29] U. N. Shenoy, Y. N. Srikant, and V. P. Bhatkar. An automatic parallelization framework for multicomputers. *Comput. Lang.*, 20(3):135–150, 1994.
- [30] J. Shirako, K. Nagasawa, K. Ishizaka, M. Obata, and H. Kasahara. Selective inline expansion for improvement of multi grain parallelism. In *Parallel and Distributed Computing and Networks*, pages 476–482, 2004.
- [31] S. Triantafyllis, M. J. Bridges, E. Raman, G. Ottoni, and D. I. August. A framework for unrestricted whole-program optimization. In *In ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation*, pages 61–71, 2006.
- [32] O. Waddell and R. K. Dybig. Fast and effective procedure inlining. In *SAS '97: Proceedings of the 4th International Symposium on Static Analysis*, pages 35–52, London, UK, 1997. Springer-Verlag.
- [33] M. J. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.
- [34] P. Wu, S. P. Midkiff, J. E. Moreira, and M. Gupta. Improving Java performance through semantic inlining. In *Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing*, Mar 1999.
- [35] P. Wu, J. E. Moreira, S. P. Midkiff, M. Gupta, and D. A. Padua. Semantic inlining - the compiler support for java in technical computing. In *PPSC*, 1999.
- [36] P. Zhao and J. N. Amaral. To inline or not to inline? enhanced inlining decisions. In *In Workshop on Languages and Compilers for Parallel Computing (LCPC)*, pages 405–419, 2003.
- [37] P. Zhao and J. N. Amaral. Ablego: a function outlining and partial inlining framework. *Softw. Pract. Exper.*, 37(5):465–491, 2007.