

Information-Flow Analysis of Design Breaks up

Rajendra kumar

Research Scholar-Department of Computer Science, GITM Lucknow, Uttar Pradesh technical University, Lucknow.
Email: rajendra888999@gmail.com

Anil Kumar

Research Scholar-Department of Computer Science, GITM Lucknow, Uttar Pradesh technical University, Lucknow.
Email: anilgangwar15@gmail.com

Namrata Dhanda

Assott. Professor-Department of Computer Science, GITM Lucknow, Uttar Pradesh technical University, Lucknow.
Email: ndhanda510@gmail.com

ABSTRACT

Traditional interpreted data-flow analysis is executed on whole plans; however, such **whole-program psychoanalysis** is not executable for large or uncompleted plans. We suggest **fragment data-flow analysis** as a substitute approach which calculates data-flow information for a particular program fragment. The psychoanalysis is parameterized by the extra information available about the rest of the program. We depict two frameworks for interracial flow-sensitive fragment psychoanalysis, the relationship amongst fragment psychoanalysis and whole-program analysis, and the necessities ensuring fragment analysis safety and feasibility. We suggest an application of fragment analysis as a second analysis phase after a cheap flow-insensitive whole-program analysis, in order to obtain better data for important program fragments. We also depict the design of two fragment analyses derived from an already existing whole-program flow- and context-sensitive pointer alias analysis for Computer program and present empirical rating of their cost and precision. Our experiments show evidence of dramatic improves precision gettable at a practical cost.

Date of Submission: February 11, 2014

Date of Acceptance: March 12, 2014

1 Introduction

Many stages of the software development cycle require information about the properties of large and complex programs. *Data-flow analysis* extracts semantic information which can be used for code optimization, program slicing, semantic change analysis, program restructuring, and code testing. In many cases, *interprocedural* data-flow analysis is needed to obtain information about program properties that depend on the interaction between different procedures. *Flow insensitive*

Analysis dismisses the governing of arguments and computes one answer for the whole program; in contrast, *flow-sensitive* depth psychology follows the control flow order of statements and computes different solutions at distinct program points. *Context-sensitive* analysis conceives (sometimes approximately) only paths along which calls and returns are decently matched, while *context-insensitive* psychoanalysis does not make this eminence.

Traditionally, interprocedural data-flow psychoanalysis is planned to examine whole programs; however, in many cases such *whole-program analysis* is impracticable. For This research was supported, in part, by NSF grants CCR-9501761 and CCR- 9804065, and Edison Design Group. Very large programs with hundreds of thousands or even millions lines of code, the time commanded to build a whole-program representation and

the space demanded to store it are prohibitive [1]. In many cases, the plans are incomplete the source code for components of the program (e.g., libraries) is not uncommitted. Empirical evidence proposes that precise interprocedural flow-sensitive analysis does not scale for large programs [18]. In approximately cases the analysis results are not needed for the whole program, but only for a relatively small part of it | for example, a maintenance task for a specific program fragment may only expect data-flow information for curriculum directs inside the break up, both before and after the sustentation change. This paper suggests an alternative approach for plan analysis. Instead of addressing the problem of calculating data-flow information for the whole program, we deal the trouble of calculating data-flow information for a particular program break up. The problem is parameterized by the extra information available about the rest of the plan. Such *fragment data-flow psychoanalysis* can avoid the troubles of the conventional whole-program psychoanalysis. For example, data can be found about fragments of very large plans for which whole-program analysis is prohibitively expensive. Likewise, analysis can be executed on fragmentizes of incomplete programs; for such programs, traditional whole-program analysis is not possible. Eventually, fragment analysis calculates data about only the "interesting portion" of the curriculum, which can be significantly smaller than the program itself. This paper is a first step in investigation the

theory and practice of fragment data-flow psychoanalysis. It only conceives flow-sensitive fragment analysis. The main parts of this work can be summarized as follows:

-We describe two frameworks for interprocedural flow-sensitive fragment analysis, derived from existing frameworks for whole-program analysis. We discuss the relationship between fragment analysis and whole-program analysis, and the requirements ensuring fragment analysis safety and feasibility.

- We suggest a coating of fragment psychoanalysis as a second analysis phase after an inexpensive flow-insensitive whole-program analysis, in order to receive better information for important program fragments. This approach can be used for plans that are too big to be examined by flow-sensitive whole program analysis, yet allow flow-insensitive whole-program analysis for example, C programs with around 100,000 lines of code [1, 18, 17].
- We depict the design of two fragment analyzes derived from a whole program flow- and context-sensitive pointer alias analysis [10] for C programs.
- We present empirical rating of the cost and precision of these two fragment pointer alias analyses. We show that the time and space costs of the analyses are practical. In about 75% of our experiments, the better of the two analyses results in a fourfold or higher precision improvement over the whole-program flow-insensitive solution.

The rest of the paper is formed as follows: Section 2 depicts models for flow-sensitive whole-program analysis. Section 3 talks about frameworks for fragment analysis and the issues demanded in their design. Section 4 depicts the whole-program pointer alias analysis from [10], and Section 5 depicts the design of two fragment pointer alias analyses. Empirical results are demonstrated in Section 6. Section 7 describes related work, and Section 8 presents our conclusions.

2 Whole-Program Data-Flow Analysis

This section demonstrates two well-known models for whole-program interprocedural flow-sensitive data-flow analysis. Without loss of generality, we will only consider analysis for forward data-flow problems [12]. Given a whole program to be analyzed, a *whole-program analysis* constructs a data-flow framework $\langle G, L, F, M, \eta \rangle$, where:

- $G = (N, E, \rho)$ is a aimed graph with node set N , edge set E and starting node $\rho \in N$ (for our purposes, G is an interprocedural control flow graph).
- $\langle L, \leq \wedge \rangle$ is a meet semi-lattice [12] with partial order \leq and meet \wedge . For simplicity, we only consider L which is finite¹ and has a top element T .

- $F \subseteq \{f : L \rightarrow L\}$ is a function space closed under composition and arbitrary meets. We assume that F is monotone [12].
- $M : N \rightarrow F$ is an assignment of transfer functions to the nodes in G (without loss of generality, we assume no edge transfer functions). The transfer function for node n will be denoted by fn .
- $\eta \in L$ is the solution at the bottom of ρ .

The program is constituted by an *interprocedural control flow graph* (ICFG) [10], which comprises control flow graphs for all processes in the program. Each process has a single *entry node* (node ρ is the entry node of the starting procedure) and a single *exit node*. Each call statement is constituted by a pair of nodes, a *call node* and a *return node*. There is an edge from the call node to the entry node of the called process; there is also an edge from the exit node of the called process to the return node in the calling process.

A path from node n_1 to node n_k is a sequence of nodes $p = (n_1, \dots, n_k)$ such that $(n_i, n_{i+1}) \in E$. Let $fp = fn_1 \circ fn_2 \circ \dots \circ fn_k$. An *accomplishable path* is a path on which every process returns to the call site which raised it [16, 10, and 13]; only such paths represent possible sequences of execution steps. A *same level realizable path* is a accomplishable path whose first and last nodes belong to the same process, and on which the number of call nodes is equal to the number of return nodes. Such paths represent sequences of execution steps during which the call stack may temporarily grow deeper, but never shallower than its original depth, before eventually returning to its original depth [13]. The set of all achievable paths from n to m will be denoted by $RP(n, m)$; the set of all same-level realizable paths from n to m will be denoted by $SLRP(n, m)$.

Definition 1. For each $n \in N$, the *meet-over-all-realizable-paths (MORP) solution* at n is defined as $MORP(n) = \wedge_{p \in CRP(\rho, n)} fp(\eta)$.

1 The results can be easily generalized for finite-height semi-lattices.

Context-Insensitive Analysis. After constructing $\langle G, L, F, M, \eta \rangle$, a whole program analysis computes a solution $S : N \rightarrow L$; the data-flow solution at the bottom of node n will be denoted by S_n . The solution is *safe* iff $S_n \leq MORP(n)$ for each node n ; a *safe analysis* computes a safe solution for each valid input program. Traditionally, a system of equations is constructed and then solved using fixed-point iteration. In the simplest case, a context-insensitive analysis constructs a system of equations of the form

$$S_p = \eta, \quad S_n = \wedge_{m \in Pred(n)} fn(Sm)$$

Where $Pred(n)$ is the set of predecessor nodes for n . The initial solution has $S_p = \eta$ and $S_n = T$ for any $n \neq \rho$; the final

solution is a fixed point of the system and is also a safe estimate of the MORP solution.

Context-Sensitive Analysis. The problem with the above approach is that data is disseminated from the exit of a process to *all* of its callers. Context sensitive psychoanalysis is used to deal this potential source of imprecision. One approach is to disseminate factors of L collectively tags which approximate the calling circumstance of the procedure. At the exit of the process, these tags are looked up in order to back-propagate data only to call sites at which the corresponding calling context existed. The "functional approach" to context sensitivity [16] uses a new framework that is the function space of functions mapping L to L . Intuitively, if the solution at node n is a map $hn : L \rightarrow L$, then for each $x \in L$ we can use $hn(x)$ to approximate $fp(x)$ for each path $p \in SLRP(e; n)$, where e is the entry node of the procedure containing n . In other words, $hn(x)$ estimates the part of the solution at n that occurs under calling context x at e . If calling context x never occurs at e , then $hn(x) = \perp$. The solution of the original problem is received as $\bigvee x \in L hn(x)$. In general, this approach requires compact representation of functions and explicit functional compositions and meets, which are usually infeasible. When L is finite, a feasible version of the analysis can be designed [16]. Figure 1 presents a simplified description of this version. $H[n,x]$ contains the current value of $hn(x)$; the worklist contains pairs (n,x) for which $H[n,x]$ has changed and has to be propagated to the successors of n . If n is a call node, at line 7 the value of $H[n,x]$ is propagated to the entry node of the called procedure. If n is an exit node, the value of $H[n,x]$ is propagated only to return nodes at whose corresponding call nodes x occurs (lines 8-10).

For distributive frameworks [12], this algorithm finishes with the MORP solution; for non-distributive monotone frameworks, it produces a safe approximation of the MORP solution [16]. When the lattice is the power set of some basic finite set D of data-flow facts (e.g., the set all potential aliases or the set of all variable definitions), the algorithm can be modified to propagate elements of D instead of elements of $2D$. For distributive frameworks, this approach produces a precise solution [13]. For non-distributive monotone frameworks, restricting the context to a singleton set necessarily introduces some approximation; the whole-program pointer alias analysis from [10] falls in this category.

input $\langle G; L; F; M; _ \rangle$; L is finite
output S : array[N] of L
declare H : array[N, L] of L ; initial values \perp
 W : list of $(n,x), n \in N, x \in L$; initially empty

```
[1]  $H[\_, \_] := \perp$ ;  $W := f(\_, \_)g$ ;
[2] while  $W \neq \emptyset$ ; do [3] remove  $(n,x)$  from  $W$ ;  $y := H[n,x]$ ;
[4] if  $n$  is not a call node or an exit node then
[5] foreach  $m \in Succ(n)$  do propagate( $m,x,fm(y)$ );
[6] if  $n$  is a call node then
```

```
[7]  $e :=$  called entry( $n$ ); propagate( $e,y,fe(y)$ );
[8] if  $n$  is an exit node then
[9] foreach  $r \in Succ(n)$  and  $l \in L$  do
[10] if  $H[\text{call node}(r),l] = x$  then propagate( $r,l,fr(y)$ );
[11] foreach  $n \in N$  do
[12]  $S[n] := \bigvee l \in L H[n,l]$ ;
[13] procedure propagate( $n,x,y$ )
[14]  $H[n,x] := H[n,x] \wedge y$ ; if  $H[n,x]$  changed then add  $(n,x)$ 
to  $W$ ;
```

Fig. 1. Worklist implementation of context-sensitive whole-program analysis

3 Fragment Data-Flow Analysis

This section depicts how interprocedural flow-sensitive whole-program psychoanalysis can be changed to obtain *fragment data-flow analysis*, which examines a program fragment instead of a whole curriculum. The structure of context-insensitive and context-sensitive fragment analysis is discussed, as well as the issues involved in the design of fragment analysis and the requirements ensuring its safety.

3.1 Fragment Analysis Structure

We assume that the analysis input admits a *program fragment* F , which is an absolute set of processes. We expect these procedures to be powerfully interrelated; otherwise, the psychoanalysis may yield information that is too imprecise. The input also contains *whole-program information* I , which represents the knowledge available about the programs to which F belongs. The whole-program information depends on the particular software development environment and the process in which fragment analysis is used; the role of I is further discussed in Sect. 3.2. We will use $PI(F)$ to denote the set of all valid whole programs that contain F and for which I is true; depending on I , $PI(F)$ can be anything from a singleton set (e.g., when the source code of the whole program is available) to an infinite set. Given F and I , a fragment analysis extracts several kinds of information, shown in Table 1. Graph $G_0 = (N_0; E_0)$ is the ICFG for the fragment and can be constructed similarly to the whole-program case, except for calls to procedures outside of F , which are not represented by any edges in G_0 . Set *BoundaryEntries* contains every entry node $e \in N_0$ which has a predecessor $c \in N_0$ in some program from $PI(F)$. Similarly, *BoundaryCalls* contains every call node $c \in N_0$ which has a successor $e \in N_0$ in some program from $PI(F)$. Set *BoundaryReturns* comprises the return nodes corresponding to call nodes from *BoundaryCalls*.

Table 1. Information extracted by a fragment analysis

Information	Description
$\langle G'; L'; F'; M'; \eta' \rangle$	Data-flow framework
<i>BoundaryEntries</i>	Entry nodes of procedures called from outside of <i>F</i>
<i>BoundaryCalls</i>	Call nodes to procedures outside of <i>F</i>
<i>BoundaryReturns</i>	Return nodes from procedures outside of <i>F</i>
$\beta_e : \text{BoundaryEntries} \rightarrow L'$	Summary information at boundary entry nodes
$\beta_c : \text{BoundaryCalls} \rightarrow F$	Summary information at boundary call nodes

3.2 Fragment Analysis Design

Designers of a specific fragment analysis have to address several important problems. One problem is to decide what kind of whole-program information *I* to use. The decision depends on the software development environment, as well as the process in which the fragment analysis is used. In this paper, we are particularly interested in a process where an inexpensive flow-insensitive whole-program

$S0\ n = \vee m2Pred(n) f0\ n(S0\ m)$ if $n = 2\text{BoundaryEntries} [\text{BoundaryReturns}$
 $S0\ n = _e(n) \wedge \vee m2Pred(n) f0\ n(S0\ m)$
 if $n2\text{BoundaryEntries}$
 $S0\ n = f0\ n(_c(m)(S0\ m))$ if $n2\text{BoundaryReturns}$ and $m =$
 call node(n)
 $S0\ _ = _0$ if $_2N$

Fig. 2. Context-insensitive fragment analysis

input $\langle G0; L0; F0; M0; _0 \rangle$; $L0$ is $_nite$
output S : array[$N0$] of $L0$
declare H : array[$N0, L0$] of $L0$; initial values >0
 W : list of (n,x) , $n2N0, x2L0$; initially empty

[1a*] **if** $_2N0$ **then**
 [1b] $H[_, _0] := _0$; add $(_, _0)$ to W ;
 [1c*] **foreach** $n2\text{BoundaryEntries}$ **do**
 [1d*] $H[n, _e(n)] := _e(n)$; add $(n, _e(n))$ to W ;
 [2] **while** $W \neq ;$ **do**
 [3] remove (n,x) from W ; $y := H[n,x]$;
 [4] **if** n is not a call node or an exit node **then**
 [5] **foreach** $m2\text{Succ}(n)$ **do** propagate($m,x, f0\ m(y)$);
 [6a*] **if** n is a call node and $n = 2\text{BoundaryCalls}$ **then**
 [7a] $e :=$ called entry(n); propagate($e,y, f0\ e(y)$);
 [6b*] **if** n is a call node and $n2\text{BoundaryCalls}$ **then**
 [7b*] $r :=$ ret node(n); propagate($r,x, f0\ r(_c(n)(y))$);
 [8] **if** n is an exit node **then**
 [9] **foreach** $r2\text{Succ}(n)$ and $l02L$ **do**
 [10] **if** $H[\text{call node}(r), l0] = x$ **then** propagate($r, l0, f0\ r(y)$);
 [11] **foreach** $n2N0$ **do**
 [12] $S[n] := \vee l02L0\ H[n, l0]$;

Fig. 3. Worklist implementation of context-sensitive fragment analysis

analysis is performed first, and then a more precise flow-sensitive fragment analysis is used on fragments for which better information is needed. This approach can be used for programs that are too big to be analyzed by flow-sensitive wholeprogram analysis, yet allow flow-insensitive whole-program analysis [for example, C programs with around 100,000 lines of code [1, 18, 17]. In this scenario, the first stage computes a whole-program flow-insensitive solution and the program call graph. The second stage uses the call graph and the flow-insensitive solution as its whole-program information *I*. The two fragment analyses in Sect. 5 are designed in this manner.

Another problem is to construct the information described in Table 1. Sets *BoundaryCalls* and *BoundaryReturns* can be determined from *F*. Set *BoundaryEntries* by definition depends on *I*. The summary information at boundary nodes also depends on *I*. When the fragment analysis follows a flow-insensitive whole-program analysis, *I* contains the call graph of the program, from which the boundary entries can be easily determined. In this case, β_e and β_c can be extracted from the whole-program flow-insensitive solution, as shown in Sect. 5.

The semi-lattice $L0$ depends mostly on *F*. However, it may be also dependent on *I* | for example, when the fragment analysis follows a flow-insensitive whole program analysis, *I* contains a whole-program solution which may need to be represented (possibly approximately) using $L0$. The fragment analysis complexity is bounded by a function of the size of $L0$; therefore, it is crucial to ensure that the size of $L0$ depends *only* on the size of *F* and not the size of *I*. This requirement guarantees that the fragment analysis will be feasible for relatively small fragments of very large programs. The fragment analyses from Sect. 5 illustrate this situation.

3.3 Fragment Analysis Safety

The fragment analyses outlined above are similar in structure to the whole program analyses from Sect. 2. In fact, we are only interested in fragment analyses that are *derived* from whole-program analyses. Consider a safe whole-program analysis which we want to modify in order to obtain a safe fragment analysis. The most important problem is to define the relationship between the semi-lattice $L0$ for a fragment *F* and the whole-program semi-lattices Lp for the programs $p2PI(F)$. For each such Lp , the designers of the fragment analysis must define an *abstraction relation* $_p\ Lp\ _L0$. This relation encodes the notion of safety and is used to prove that the analysis is safe;

it is never explicitly constructed or used during the analysis. If $(x; x0) \in _p$, we will write $_p(x; x0)$.

Intuitively, the abstraction relation $_p$ defines the relationship between the "knowledge" represented by values from Lp and the "knowledge" represented by values from $L0$. If $_p(x; x0)$, the knowledge associated with $x0$ "safely abstracts" the knowledge associated with x ; thus, $_p$ is similar in nature to the abstraction relations used for abstract interpretation [19, 6]. The choice of $_p$ depends both on the original whole-program analysis and the intended clients of the fragment analysis solution. Sect. 5 presents an example of one such choice.

Definition 2. A solution produced by a fragment analysis for an input pair $(F; I)$ is **safe** if $_p(\text{MORPp}(n); S0\ n)$ for each $p2PI(F)$ and each $n2N0$, where $S0\ n$ is the fragment analysis solution at n and $\text{MORPp}(n)$ is the MORP solution at n in p . A **safe fragment analysis** yields a safe solution for each valid input pair $(F; I)$.

With this definition in mind, we present a set of sufficient requirements that ensures the safety of the fragment analysis. Intuitively, the first requirement ensures that a safe approximation in $L0$, as defined by the partial order $_$, is also a safe abstraction according to $_p$. The second requirement ensures that if $x0 \in L0$ safely abstracts the effects of each realizable path ending at a given node n , it can be used to safely abstract the MORP solution at n . Formally, for any $p2PI(F)$ and its $_p$, any $x; y \in Lp$ and any $x0; y0 \in L0$, the following must be true:

Property 1: if $_p(x; x0)$ and $y0 _ x0$, then $_p(x; y0)$

Property 2: if $_p(x; x0)$ and $_p(y; x0)$, then $_p(x \wedge y; x0)$

The next requirement ensures that the transfer functions in the fragment analysis safely abstract the transfer functions in the whole-program analysis. Let $fn;p$ be the transfer function for $n \in N0$ in the whole-program analysis for any $p2PI(F)$. For any $n2N0$, $x \in Lp$ and $x0 \in L0$, the following must be true:

Property 3: if $_p(x; x0)$, then $_p(fn;p(x); f0\ n(x0))$

If F contains the starting procedure of the program, then $_2\ N0$ and the fragment analysis solution at the bottom of $_$ is $_0$. Let $_p$ be the whole-program solution at the bottom of $_$ for any $p2PI(F)$. Then the following must be true:

Property 4: $_p(_p; _0)$

The next requirement ensures that for each boundary entry node e , the summary value $_e(e)$ safely abstracts the effects of each realizable path that reaches e from outside of F . For any $p2PI(F)$, let $RPout\ p(_ ; e)$ be the set of all realizable paths $q = (_ ; : : ; c; e)$ in p such that $c \in N0$; recall that fq

is the composition of the transfer functions for the nodes in q . Then the following must be true:

Property 5: $\delta q2RPout\ p(_ ; e) : _p(fq(_p); _e(e))$

The last requirement ensures that the summary function $_c(c)$ for each boundary call node c safely abstracts the effects of each same-level realizable path from the entry to the exit of the called procedure. Consider any $p2PI(F)$ in which c has a successor entry node $e \in N0$ and let t be the exit node corresponding to e . Let $SLRPp(e; t)$ be the set of all same-level realizable paths in p from e to t . Intuitively, each realizable path ending at t has a suffix which is a same-level realizable path; thus, $_c(c)$ should safely abstract each path from $SLRPp(e; t)$. Formally, for any $x \in Lp$ and $x0 \in L0$, the following must be true:

Property 6: $\delta q2SLRPp(e; t)$: if $_p(x; x0)$, then $_p(fq(x); _c(c)(x0))$

If the above requirements are satisfied, the context-insensitive fragment analysis derived from a safe context-insensitive whole-program analysis is safe, according to Definition 2. The proof considers a fixed-point solution of the system in Fig. 2. It can be shown that for each $p \in PI(F)$, each $n \in N0$ and each realizable path $q = (_p; : : ; n)$, it is true that $_p(fq(_p); S0\ n)$. Each such q is the concatenation of two realizable paths $r0$ and $r00$. Path $r0$ starts from $_p$ and lies entirely outside of F ; if $_p \in N0$, $r0$ is empty. Path $r00 = (e; : : ; n)$ has $e \in N0$ and $n \in N0$ and is the *fragment suffix* of q ; it may enter and leave F arbitrarily. The proof is by induction on the length of the fragment suffix of q .

Similarly, the context-sensitive fragment analysis derived from a safe context sensitive whole-program analysis is safe. It is enough to show that for each $p2PI(F)$, each $n \in N0$ and each realizable path $q = (_p; : : ; n)$, there exists a value $l0 \in L0$ such that $_p(fq(_p); H[n; l0])$. Again, this can be proven by induction on the length of the fragment suffix of q .

4 Whole-Program Pointer Alias Analysis

This section presents a simplified high-level description of the Landi-Ryder whole-program pointer alias analysis [10] for the C programming language. The analysis considers a set of names that can be described by the grammar in Fig. 4. Each of the names corresponds to one or more run-time memory locations. An *alias pair* (or simply an *alias*) is a pair of names that potentially represent the same memory location.

```
<Name> ::= <SimpleName> j <Deref> j <ArrayElem>
          | <FieldAccess> j <K-Limited>
<SimpleName> ::= identifier /* variable name */
                | heapn /* heap location */
```

```

<Deref> ::= *(<Name>+?) /* dereference */
<ArrayElem> ::= <Name>[?] /* array element */
<Fld> ::= identifier /* field name */
<FieldAccess> ::= <Name>.<Fld> /* field of a structure */
<K-Limited> ::= <Name># /* k-limited name */
    
```

Fig. 4. Grammar for names

If a program point n is a call to malloc or a similar function, name $heapn$ represents the set of all heap memory locations created at this point during execution. If there are recursive data structures (e.g., linked lists), the number of names is potentially infinite. The analysis limits the number of dereferences in a name to a given constant k ; any name with more than k dereferences is represented by a k -limited name. For example, for $k = 1$, name $_((_p):f):f$ is represented by the k -limited name $_((_p):f)\#$.

A *simple name* is a name generated by the SimpleName nonterminal in the grammar. The *root name* of a name n is the simple name used when n is generated by the grammar; for example, the root name of $_((_p):f)$ is p . Name n is a *fixed location* if it does not contain any dereferences. Name $_((_p):f)$ is not a fixed location, while $s[?]:f$ is a fixed location.

The lattice of the analysis is the power set of the set of all pairs of names; the meet operator is set union and the partial order is the "is-superset-of" relation. The analysis is flow- and context-sensitive, and conceptually follows the Sharir-Pnueli algorithm described in Sect. 2. However, since the lattice is a power set, the algorithm can be modified to propagate single aliases instead of sets of aliases | the worklist contains triples (n, RA, PA) , where n is a node, RA is a reaching alias that is part of the solution at the entry of the procedure to which n belongs, and PA is a possible alias at the bottom of n . Restricting the reaching alias set to a single alias introduces some approximation; therefore, the actual Landi-Ryder algorithm can be viewed as an approximation algorithm solving the more precise Sharir-Pnueli formulation of the problem.

The fragment analyses in Sect. 5 are derived from a modified version of the Landi-Ryder analysis. This version considers only aliases containing a fixed location; aliases with two non-fixed locations (e.g., an alias between $_((_p):f)$ and $_((_q):g)$) are ignored during propagation. It can be shown that this is safe, as long as the program does not contain assignments in the form $*p=\&x$; if such assignments exist, they can be removed by introducing intermediate temporary variables | for example, $t=\&x$; $*p=t$. An assignment whose left-hand side is a non-fixed location (i.e., *through-deref assignment*) potentially modifies many fixed locations. A safe solution for the modified analysis is sufficient to estimate all fixed locations possibly modified by through-deref assignments. We refer to this process as *resolution* of through-deref assignments.

5 Fragment Pointer Alias Analyses

This section describes two fragment pointer alias analyses | basic analysis $A1$ and extended analysis $A2$ | derived from the modified whole-program analysis in Sect. 4. They are used for resolution of through-deref assignments. Their design is based on a process in which flow-insensitive whole-program pointer alias analysis is performed $_rst$, and then more precise flow-sensitive fragment pointer alias analysis is used on fragments for which better information is needed. As described in Sect. 3.2, in this case the whole-program information I contains a whole-program flow-insensitive solution and the program call graph.

Analysis Input. The flow-insensitive solution SFI is obtained using a wholeprogram flow- and context-insensitive analysis similar to the one in [21]. Intuitively, the names in the program are partitioned into equivalence classes; if two names can be aliases, they belong to the same class in the final solution. Every name starts in its own equivalence class and then classes are joined as possible aliases are discovered. For example, statement $p=\&x$ causes the equivalence classes of $_p$ and x to merge. Overall, the analysis is similar to other flow- and context-insensitive analyses with almost-linear cost [17, 15]. Then SFI is used to resolve calls through function pointers and to produce a safe approximation of the program call graph. The sets of boundary entry, call and return nodes can be easily determined from this graph.

The basic analysis $A1$ takes as input SFI and the program call graph, as well as the source code for the analyzed fragment F . The extended analysis $A2$ requires as additional input the source code for all procedures directly or indirectly called by procedures in F . These additional procedures together with the procedures from F form the *extended fragment Fext*. Including all transitively called procedures allows $A2$ to estimate better the effects of calls to procedures outside of F ; the tradeoffs of this approach are further discussed in Sect. 8.

Analysis Lattices. The whole-program lattice is based on the set of names generated by the grammar in Fig. 4. Each such name can be classified as either relevant or irrelevant. A *relevant name* has a root name with a syntactic occurrence in the fragment; all other names are irrelevant. Since the fragment analysis solution is used to resolve through-deref assignments, only aliases that contain a relevant non-fixed location are useful. If the new lattice $L0$ contains all such aliases, its size still potentially depends on the number of fixed locations in the whole program. This problem can be solved by using special *placeholder variables* to represent sets of related irrelevant fixed locations; this approach is similar to the use of representative data-flow values in other analyses [11, 10, 7, 20, 4].

Consider an equivalence class Ei 2 SFI that contains at least one relevant name. If Ei contains irrelevant fixed locations,

a placeholder variable phi is used to represent them. The aliases from $L0$ fall in two categories: (1) a pair of relevant names, exactly one of which is a fixed location, and (2) a relevant non-fixed location and a placeholder variable. Clearly, the number of relevant names and the number of placeholder variables depend only on the size of the fragment; thus, the size of $L0$ is independent of the size of the whole program.

In the final solution, an alias with two relevant names represents itself, while an alias with a placeholder phi represents a set of aliases, one for each irrelevant fixed location represented by phi . This can be formalized by defining an abstraction relation $_p _Lp _L0$ for each $p2PI$ (F). For each pair of sets $S _Lp$ and $S0 _L0$, we have $_p(S; S0)$ i_ each alias from S is safely abstracted by $S0$. Let x be the non-fixed location in alias $(x,y)2S$; then $_p$ is defined as follows:

- If x is an irrelevant name, then (x,y) is safely abstracted by $S0$
- If x and y are relevant and $(x,y)2S0$, then (x,y) is safely abstracted by $S0$
- If x is relevant and y is irrelevant, Ei is y 's equivalence class, and $(x,phi)2S0$, then (x,y) is safely abstracted by $S0$

Intuitively, the above definition says that aliases involving an irrelevant non- fixed location can be safely ignored. It also says that irrelevant fixed locations from the same equivalence class have equivalent behavior and need not be considered individually. The safety implications of this definition are discussed later. Note that if a fragment alias solution is safe according to the above definition, it can be used to safely resolve through-deref assignments.

Summaries at Boundary Nodes. Based on SFI , $A1$ and $A2$ construct a set of aliases $_L0$. Each equivalence class Ei is examined and for each pair of names (x,y) of a non-fixed location x and a fixed location y from the class, the following is done: first, if x is an irrelevant name, the pair is ignored. Otherwise, if y is a relevant name, (x,y) is added to $_$. Otherwise, if y is an irrelevant name, (x,phi) is added to $_$. Clearly, (x,y) is safely abstracted by $_$ according to the definition of $_p$ given above.

The basic analysis $A1$ defines $_e(e) = _$ for each boundary entry node e and $_c(c)(x) = _$ for each $x _L0$ and each boundary call node c . This essentially means that SFI is used to approximate the solutions at boundary entry nodes and boundary return nodes. For the extended analysis $A2$, there are no boundary call nodes. For each boundary entry node e that belongs to the original fragment F , the analysis defines $_e(e) = _$. If e is part of the extended fragment $Fext$, but is not part of the original fragment F , the analysis defines $_e(e) = \theta$.

Analysis Safety. To show the safety of $A1$, it is enough to show that the requirements from Sect. 3.3 are satisfied. Properties 1 and 2 are trivially satisfied. Since both $_p$ and $_0$ are the empty set, Property 4 is also true.

Showing that Property 3 is true requires careful examination of the transfer functions from [10]. The formal proof is based on two key observations. The first one is that aliases involving an irrelevant non-fixed location are only propagated through the nodes in the fragment, without actually creating any new aliases; therefore, they can be safely ignored. The second observation is that aliases with the same no fixed location and with different irrelevant fixed locations have equivalent behavior. For example, alias $(_p,x)$ at the top of statement $q=p$; results in $(_p,x)$ and $(_q,x)$ at the bottom of the statement. Similarly, $(_p,y)$ results in $(_p,y)$ and $(_q,y)$. If both x and y are represented by a placeholder phi , alias $(_p,phi)$ results in $(_p,phi)$ and $(_q,phi)$, which satisfies Property 3.

The set $_$ described above is extracted from the whole-programflow-insensitive alias solution, which is safe. Therefore, $_$ safely abstracts any aliases that could be true at a node in the program; thus, Properties 5 and 6 are true. Since all requirements are satisfied, $A1$ is safe. Similarly to the whole-program case, the actual implementation propagates single aliases instead of sets of aliases. As a result, the reaching alias set is restricted to a single alias and therefore some approximation is introduced. Of course, the resulting solution is still safe.

For the extended analysis $A2$, it is *not true* that the solution is safe at each node in the extended fragment. For example, consider the entry node of a procedure that was not in the original fragment F . If this procedure is called from outside of the extended fragment $Fext$, aliases could be propagated along this call edge during the whole-program analysis, but would be missing in the fragment analysis. However, it can be proven that for each node in the original fragment F , the solution is safe. The proof is very similar to the one outlined in Sect. 3.3, and still requires that Properties 1 through 4 are true. For each realizable path q starting at $_$ and ending at a node in F , its fragment suffix is the subpath starting at the first node in q that belongs to F . The proof is by induction on the length of the fragment suffix; the base case of the induction depends on the fact that $_e(e) = _$ for boundary entry nodes in F , and therefore the solution at such nodes is safe.

6 Empirical Results

Our implementation uses the PROLANGS Analysis Framework2 (version 1.0), which incorporates the Edison Design Group front end for C/C++. The results were gathered on a *Sun Sparc-20* with 352 MB of memory. The implementation analyzes a reduced version of C that

excludes signals, setjmp, and longjmp, but allows function pointers, type casting and union types.

Table 2 describes the C programs used in our experiments. It shows the program size in lines of code and number of ICFG nodes, the number of procedures, the total number of assignments, and the number of through-deref assignments.

Table 2. Analyzed Programs

Fragment	ICFG Nodes	% Total	Procs	Boundary		Assignments	
				Entries	Calls	All	Deref
zip.1	1351	21.9	28	5	17	776	59
zip.2	429	7.0	9	5	19	255	50
sc.1	1238	18.5	30	3	30	609	8
sc.2	793	11.9	13	3	10	459	23
larn.1	345	2.9	4	4	35	188	46
larn.2	420	3.5	11	9	44	216	4
espresso.1	440	2.9	6	2	40	234	38
espresso.2	963	6.3	19	5	113	461	53
tsl.1	355	2.3	13	4	33	175	15
tsl.2	1004	6.5	29	7	134	459	17
moria.1	2678	13.2	43	14	348	1634	392
moria.2	1221	6.0	27	7	149	644	49

For each of the data programs, we extracted by hand subsets of related procedures that formed a cohesive fragment. Significant effort was put in determining realistic fragments | the program source code was thoroughly examined, the program call graph was used to obtain better understanding of the calling relationships in the program, and the documentation (both external and inside the source code) was carefully analyzed. For each program, two fragments were extracted. For example, for zip one of the fragments consisted of all procedures implementing the implosion algorithm. For espresso, one of the fragments consisted of the procedures for reducing the cubes of a boolean function. For sc, one of the fragments consisted of the procedures used to evaluate expressions in a spreadsheet. Some characteristics of the fragments are given in Table 3.

For each fragment, three experiments were performed; the results are shown in Table 4. In the first experiment, the solution from the whole-program flowinsensitive (FI) analysis was used at each through-deref assignment to determine the number of simple names possibly modified by the assignment. For a fixed location that was not a simple name, a modification of its root simple name was counted (e.g., a modification of *s:f* was counted as a modification of *s*). Then

Table 3. Analyzed Fragments

Program	LOC	ICFG Nodes	Procs	Assignments		Program	LOC	ICFG Nodes	Procs	Assignments	
				All	Deref					All	Deref
zip	8177	6172	122	3443	582	espresso	14910	15339	372	7822	1322
sc	8530	6678	160	3440	159	tsl	16053	15469	471	7249	507
larn	10014	12063	298	6021	389	moria	25292	20213	458	10557	1358

Table 4. Precision Comparison

Fragment	WholePrgm FI	Basic FFS	Percent	Extended FFS	Percent
zip.1	15.78	12.98	82.3	12.98	82.3
zip.2	3.88	1.02	26.3	1.02	26.3
sc.1	9.50	9.00	94.7	1.00	10.5
sc.2	13.14	3.29	25.0	3.29	25.0
larn.1	17.00	14.57	85.7	1.00	5.9
larn.2	9.00	9.00	100.0	1.00	11.1
espresso.1	107.10	107.10	100.0	34.00	31.7
espresso.2	57.55	57.55	100.0	35.11	61.0
tsl.1	41.87	17.07	40.8	1.00	2.4
tsl.2	41.88	24.82	59.3	1.00	2.4
moria.1	132.80	1.46	1.1	1.46	1.1
moria.2	31.04	3.55	11.4	3.55	11.4

the average across all through-deref assignments in the fragment was taken. In the second experiment, the FI analysis was followed by the basic fragment flowsensitive (FFS) analysis. Again, for each through-deref assignment the number of simple names possibly modified was determined; placeholder variables were expanded to determine the actual simple names modified. Then the average across all through-deref assignments was taken. Each of these averages is shown as an absolute value and as a percent of the FI average. The third experiment used the extended fragment analysis instead of the basic fragment analysis.

Overall, the results show that the precision of the extended analysis is very good. In particular, for seven fragments it produces averages very close to 1, which is the lower bound. The averages are bigger than 4 for only three fragments; all three take as input a pointer to an external data structure, and a large number of the through-deref assignments in the fragment are through this pointer. The pointer itself is not modified, and each modification through it resolves to the same number of simple names as in the FI solution.

The performance of the basic analysis is less satisfactory. For five fragments, it achieves the same precision as the extended analysis. For the remaining fragments, in five cases the solution is close to the FI solution. The main reason is that precision gains from flow-sensitivity are lost at calls to procedures outside of the fragment.

Table 5 shows the running times of the analyses in minutes and seconds. The last two columns give the used space in kilobytes. The results show that the cost of the fragment analyses is acceptable, in terms of both space and time.

7 Related Work

In Harrold and Rothermel's separate pointer alias analysis [8], a software module is analyzed separately and later linked with other modules. The analysis is based

Table 5. Analysis Time and Space

Fragment	WholePr gm FI	Basi c FFS	Extend ed FFS	Basic Space	Extende d Space
zip.1	0:02	1:01	1:35	2544	3784
zip.2	0:02	0:18	0:18	2064	2064
sc.1	0:02	0:18	0:18	2504	2504
sc.2	0:02	0:25	0:33	2664	2824
larn.1	0:03	0:22	0:27	3760	6560
larn.2	0:03	0:20	0:28	3680	6556
espresso.1	0:07	0:52	1:58	5504	17400
espresso.2	0:07	0:56	3:07	5504	26904
tsl.1	0:08	0:33	0:43	5776	8672
tsl.2	0:08	1:15	1:25	12480	18648
moria.1	0:09	1:22	1:29	9504	10464
moria.2	0:09	1:39	1:37	7608	17440

on the whole-program analysis from [10]; it simulates the aliasing effects that are possible under all calling contexts for the module. Placeholder variables are used to represent sets of variables that are not explicitly referenced in the module, similarly to the placeholder variables in our fragment pointer alias analyses. Aliases are assigned extra tags describing the module calling context.

There are several differences between our work and [8]. First, we have designed a general framework for fragment analysis and emphasized the importance of the theoretical requirements that ensure analysis safety and feasibility. Second, the intended application of our fragment pointer alias analyses is to improve the information about a part of the program after an inexpensive whole-program analysis; the application in [8] is separate analysis of single-entry modules. Finally, [8] does not present empirical evaluation of the performance of the analysis; we believe that their approach may have scalability problems.

Reference [11] presents an analysis that decomposes the program into regions in which several local problems are solved. Representative values are used for actual data-flow information that is external to the region; our placeholder variables are similar to these representative values. Other similar mechanisms are the non-visible names from [10, 7], extended parameters from [20], and unknown initial values from [4]. We use an abstraction relation to capture the correspondence between the representative and the actual data-flow information; this is similar to the use of abstraction relations in the field of abstract interpretation [19, 6].

The work in [3] also addresses the analysis of program fragments and uses the notion of representative data-flow information for external data-flow values. However, in [3] the specific fragments are libraries with no boundary calls, the analysis computes def-use associations in object-oriented languages with exceptions, and there is no assumption of available whole-program information.

Cardelli [2] considers separate type checking and compilation of program fragments. He proposes a theoretical framework in which program fragments are separately compiled in the context of some information about missing fragments, and later can be safely linked together.

Model checking is a technique for verifying properties of finite-state systems; the desired properties are specified using temporal-logic formulae. Modular model checking verifies properties of system modules, under some assumptions about the environment with which the module interacts [9]. These assumptions play an analogous role to that of the whole-program information in fragment analysis. Further discussion of the relationship between data-flow analysis and model checking is given in [14].

There is some similarity in the problem addressed by our work and that in [5], which presents an analysis of modular logic programs for which a compositional semantics is defined. Each module can be analyzed using an abstract interpretation of the semantics. The analysis results for separate modules can be composed to yield results for the whole program, or alternatively, the results for one module can be used during the analysis of another module.

8 Conclusions

This paper is a first step in investigating the theory and practice of fragment data-flow analysis. It proposes fragment analysis as an alternative to traditional whole-program analysis. The theoretical issues involved in the design of safe and feasible flow-sensitive fragment analysis are discussed. One possible application of fragment analysis is to be used after a whole-program flow-insensitive

analysis in order to improve the precision for an interesting portion of the program. This paper presents one such example, in which better information about modifications through pointer dereference is obtained by performing flow- and context-sensitive fragment pointer alias analysis. The design of two such analyses is described, and empirical results evaluating their cost and precision are presented.

The empirical results show that the extended analysis presented in Sect. 5 can achieve significant precision benefits at a practical cost. The performance of the basic analysis is less satisfactory, even though in about half of the cases it achieves the precision of the extended analysis. Clearly, in some cases the whole-program flow-insensitive solution is not a precise enough estimate of the effects of calls to external procedures. The extended analysis solves this problem for the fragments used in our experiments. We expect this approach to work well in cases when the extended fragment is not much bigger than the original fragment. One typical example would be a fragment with calls only to relatively small external procedures which provide simple services; in our experience, this is a common situation. However, in some cases the extended fragment may contain a prohibitively large part of the program; furthermore, the source code for some procedures may not be available. We are currently investigating different solutions to these problems.

Acknowledgments. We would like to thank Thomas Marlowe, Matthew Arnold, and the anonymous reviewers for their valuable comments.

References

- [1]. D. Atkinson and W. Griswold. Effective whole-program analysis in the presence of pointers. In *Proc. Symp. on the Foundations of Software Engineering*, pages 46{55, 1998.
- [2]. L. Cardelli. Program fragments, linking and modularization. In *Proc. Symp. On Principles of Prog. Lang.*, pages 266{277, 1997.
- [3]. R. Chatterjee and B. G. Ryder. Data-flow-based testing of object-oriented libraries. Technical Report DCS-TR-382, Rutgers University, 1999.
- [4]. R. Chatterjee, B. G. Ryder, and W. Landi. Relevant context inference. In *Proc. Symp. on Principles of Prog. Lang.*, pages 133{146, 1999.
- [5]. M. Codish, S. Debray, and R. Giacobazzi. Compositional analysis of modular logic programs. In *Proc. Symp. on Principles of Prog. Lang.*, pages 451{464, 1993.
- [6]. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixed points. In *Proc. Symp. on Principles of Prog. Lang.*, pages 238{252, 1977.
- [7]. M. Emami, R. Ghiya, and L. J. Hendren. Context-sensitive interprocedural points-to analysis in the presence of function pointers. In *Proc. Conf. on Prog. Lang. Design and Implementation*, pages 242{257, 1994.
- [8]. M. J. Harrold and G. Rothermel. Separate computation of alias information for reuse. *IEEE Transactions on Software Engineering*, 22(7):442{460, July 1996.
- [9]. O. Kupferman and M. Y. Vardi. Modular model checking. In *Proc. Symp. on Compositionality*, LNCS 1536, pages 381{401, 1997.
- [10]. W. Landi and B. G. Ryder. A safe approximation algorithm for interprocedural pointer aliasing. In *Proc. Conf. on Prog. Lang. Design and Implementation*, pages 235{248, 1992.
- [11]. T. Marlowe and B. G. Ryder. An efficient hybrid algorithm for incremental data flow analysis. In *Proc. Symp. on Principles of Prog. Lang.*, pages 184{196, 1990.
- [12]. T. Marlowe and B. G. Ryder. Properties of data flow frameworks: A unified model. *Acta Informatica*, 28:121{163, 1990.
- [13]. T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Proc. Symp. on Principles of Prog. Lang.*, pages 49{61, 1995.
- [14]. D. Schmidt. Data flow analysis is model checking of abstract interpretations. In *Proc. Symp. on Principles of Prog. Lang.*, pages 38{48, 1998.
- [15]. M. Shapiro and S. Horwitz. Fast and accurate flow-insensitive points-to analysis. In *Proc. Symp. on Principles of Prog. Lang.*, pages 1{14, 1997.
- [16]. M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In S. Muchnick and N. Jones, editors, *Program Flow Analysis: Theory and Applications*, pages 189{234. Prentice Hall, 1981.
- [17]. B. Steensgaard. Points-to analysis in almost linear time. In *Proc. Symp. on Principles of Prog. Lang.*, pages 32{41, 1996.
- [18]. P. Stocks, B.G. Ryder, W. Landi, and S. Zhang. Comparing flow- and contextsensitivity on the modification side-effects problem. In *Proc. International Symposium on Software Testing and Analysis*, pages 21{31, 1998.
- [19]. R. Wilhelm and D. Maurer. *Compiler Design*. Addison-Wesley, 1995.
- [20]. R. Wilson and M. Lam. Efficient context-sensitive pointer analysis for C programs. In *Proc. Conf. on Prog. Lang. Design and Implementation*, pages 1{12, 1995.
- [21]. S. Zhang, B. G. Ryder, and W. Landi. Program decomposition for pointer aliasing:
- [22]. A step towards practical analyses. In *Proc. Symp. on the Foundations of Software Engineering*, pages 81{92, 1996.