

CS 526

Advanced


Compiler

Construction

<http://misailo.cs.illinois.edu/courses/cs526>

POINTER ANALYSIS

The slides adapted from Vikram Adve



Course

So far:


- Dataflow analysis (examples and theory)
- Dependency analysis
- SSA (sparse dataflow analysis via def-use chains)

Coming up next:

- Pointer analysis (generalize the dependence relationship)
- Interprocedural analysis (how to analyze function calls?)
- Abstract interpretation (generalize dataflow analysis)
- Fun topics (probabilistic, autotuning, VMs ...)

POINTER ANALYSIS

The slides adapted from Vikram Adve



Pointer Analysis

Pointer and Alias Analysis are fundamental to reasoning about heap manipulating programs (pretty much all programs today).

- **Pointer Analysis:**
 - What objects does each pointer points to?
 - Also called points-to analysis
- **Alias Analysis:**
 - Can two pointers point to the same location?
 - Client of pointer analysis

Example

X = 1

P = &X

*P = 2

return X

// What is the value of X?

Aliases

Consider references r_1 or r_2 ,

- may be of the form “ x ” or “ $*p$ ” “ $**p$ ”, “ $(*p) \rightarrow q \rightarrow i$ ”...
- We assume C notation for dereferencing pointers ($*$, \rightarrow)

Alias: r_1 and r_2 are aliased if the memory locations accessed by r_1 and r_2 overlap.

Alias Relation: A set of ordered pairs $\{(r_i, r_j)\}$ denoting aliases that *may hold* at a particular point in a program.

- Sometimes called a *may-alias* relation.

May or Must: A kind of aliasing if it happens optionally or always

- May: e.g., depending on the control flow: `if (b) { p = &q; }`
- Must: determined that they always represent aliases

Aliases

We look at the language that extends the simple expressions with the additional pointer-like structures:

$p := \&x$
| $p := q$
| $*p := q$
| $p := *q$

Consider references r_1 or r_2 ,

- may be of the form “ x ” or “ $*p$ ” “ $**p$ ”, “ $(*p) \rightarrow q \rightarrow i$ ”...
- We assume C notation for dereferencing pointers ($*$, \rightarrow)

Example

$X = 1$

$P = \&X$

$Q = P$

$*P = 2$

Alias: r1 and r2 are aliased if the memory locations accessed by r1 and r2 overlap.

Example

Aliasing pairs

X = 1

P = &X

Q = P

*P = 2

// (*P, X)

// { (*P, X), (*Q, X), (*P, *Q) }

Points-To Facts

Points-to Pair: pair $(r1, r2)$ denoting that one of the memory locations of $r1$ may hold the address of one of the memory locations of $r2$.

- Also written: $r1 \rightarrow r2$, means “ $r1$ points to $r2$ ”.

Points-to Set: $\{(r_i, r_j)\}$: A set of points-to pairs that may hold at a particular point in a program.

Points-To Graph: A directed graph where

- **Nodes** represents one or more memory objects;
- Each *Edge* $p \rightarrow q$ means some object in the node p may hold a pointer to some object in the node q .

Example

Points-to Pair: pair $(r1, r2)$ denoting that one of the memory locations of $r1$ An ordered may hold the address of one of the memory locations of $r2$.

$X = 1$

$P = \&X$

$Q = P$

$*P = 2$

Points-to pairs

// (P, X)

// $\{ (P, X), (Q, X) \}$

Example

Points-to Pair: pair $(r1, r2)$ denoting that one of the memory locations of $r1$ An ordered may hold the address of one of the memory locations of $r2$.

$X = I$

$P = \&X$

$Q = P$

$R = Q$

Points-to pairs

// (P, X)

// $\{ (P, X), (Q, X) \}$

// $\{ (P, X), (Q, X), (R, X) \}$

“Short notation”: vs the long one that would list all the aliases.

Challenges of Points-To Analysis

- **Pointers to pointers**, which can occur in many ways: take address of pointer; pointer to structure containing pointer; pass a pointer to a procedure by reference
- **Aggregate objects**: structures and arrays containing pointers
- **Recursive data structures** (lists, trees, graphs, etc.) closely related problem: anonymous heap locations
- **Control-flow**: analyzing different data paths
- **Interprocedural**: a location is often accessed from multiple functions; a common pattern (e.g., pass by reference)
- Compile-time cost
 - Number of variables, $|V|$, can be large
 - Number of alias pairs at a point can be $O(|V|^2)$

Common Simplifying Assumptions

Aggregate objects: arrays (and perhaps structures) containing pointers

Simple solution: Treat as a single big object!

- Commonplace for arrays.
- Not a good choice for structures?
 - See Intel Paper (Ghiya, Lavery & Sehr, PLDI 2001)
- Pointer arithmetic is only legal for traversing an array:
 $q = p \pm i$ and $q = \&p[i]$ are handled the same as $q = p$

Common Simplifying Assumptions

Recursive data structures (lists, trees, graphs, ...)

Solution: Compute aliases, not “shape”

- Don't prove something is a linked-list or a binary tree (leave that for *shape analysis*)
- ***k-limiting***: only track k or fewer levels of dereferencing
- Use simplified naming schemes for heap objects (later slide)

Common Simplifying Assumptions

Control-flow: analyzing different data paths blows up the analysis time/space

Solution(?): Could ignore the issue and compute a single common result for any path!

No consensus on this issue! (Will discuss later)

Naming Schemes for Heap Objects

The Naming Problem: Example I

```
int main() {  
    // Two distinct objects  
    T* p = create(n);  
    T* q = create(m);  
}
```

```
T* create(int num) {  
    // Many objects allocated here  
    return new T(num);  
}
```

Q. Should we try to distinguish the objects created in main()?

Naming Schemes for Heap Objects

The Naming Problem: Example 2

```
T* makelist(int len) {  
    T* newObj = new T; // Many distinct objects  
                      // allocated here  
    newObj->next = (--len == 0)? NULL :  
                  makelist(len);  
}
```

Q. Can we distinguish the objects created in makelist()?

Possible Naming Abstractions

H_0 : One name for the entire heap

H_T : One name per type T (for type-safe languages)

H_L : One name per heap allocation site L (line number)

H_C : One name per (acyclic) call path C (“*cloning*”)

H_F : One name per immediate caller F or call-site
 (“*one-level cloning*”)

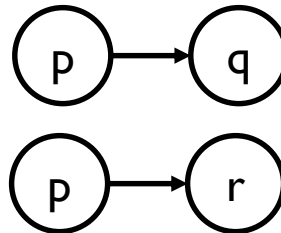
Flow-Sensitivity of Analysis

Def. A *flow-sensitive analysis* is one that computes a distinct result for each program point. A *flow-insensitive analysis* generally computes a single result for an entire procedure or an entire program.

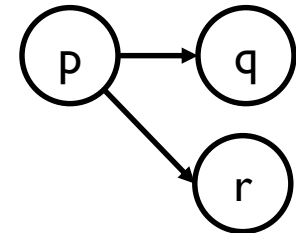
A flow-insensitive algorithm effectively ignores the order of statements!

```
int f(T q, T r){  
    T* p;  
    ...  
    p = &q;  
    ...  
    p = &r;  
}
```

Flow Sensitive



Flow Insensitive



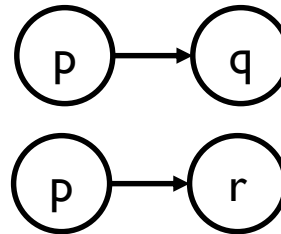
Flow-Sensitivity of Analysis

Def. A *flow-sensitive analysis* is one that computes a distinct result for each program point. A *flow-insensitive analysis* generally computes a single result for an entire procedure or an entire program.

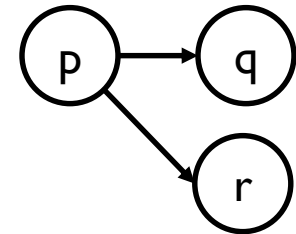
A flow-insensitive algorithm effectively ignores the order of statements!

```
int f(T q, T r){  
    T* p;  
    if (...)  
        p = &q;  
    else  
        p = &r;  
}
```

Flow Sensitive



Flow Insensitive



Flow-Sensitivity of Analysis

Pointer Analysis

- **Flow-sensitive** : At program point n , compute alias pairs $\langle a, b \rangle$ that may hold at n for some path from program entry to n .
- **Flow-insensitive** : Compute all alias pairs $\langle a, b \rangle$ such that a may be aliased to b at *some* point in a program (or function).

Important special cases

- Local scalar variables: SSA form gives flow-sensitivity
- Malloc or new: Allocates “fresh” memory, i.e., no aliases
- Read-only fields: e.g., array length

Realizable Paths

Definition: Realizable Path

A program path is realizable iff every procedure call on the path returns control to the point where it was called (or to a legal exception handler or program exit)

Whole-program Control Flow Graph?

Conceptually extend CFG to span whole program:

- split a call node in CFG into two nodes: CALL and RETURN
- add edge from CALL to ENTRY node of each callee
- add edge from EXIT node of each callee to RETURN

Problem: This produces many unrealizable paths

Focusing only on **realizable paths** requires **context-sensitive analysis**

Context-Sensitivity of Analysis

Def. A context-sensitive interprocedural analysis computes results that may hold only for realizable paths through the program.

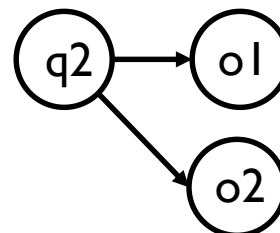
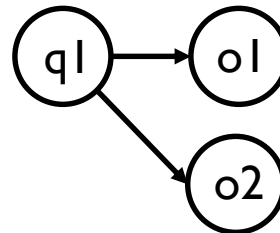
Otherwise, the analysis is context-insensitive.

```
T* identity(T* p) {  
    return p;  
}
```

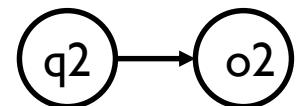
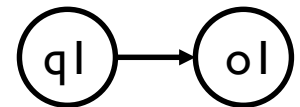
```
void f1() {  
    T* p1 = new T; // Object o1  
    T* q1 = identity(p1);  
}
```

```
void f2() {  
    T* p2 = new T; // Object o2  
    T* q2 = identity(p1);  
}
```

Context Insensitive



Context Sensitive



Context-Sensitivity of Analysis

Pointer Analysis

Apply the definitions directly using points-to pairs $\langle a, b \rangle$.

But important variations exist:

- Heap cloning vs. no cloning: Cloning gives greater context-sensitivity
- Bottom-up vs. top-down: Does final result for a procedure include only “realizable” behavior from all contexts?
- Handling of recursive functions: Does analysis retain context-sensitivity within SCCs in the call graph?

Object Sensitivity: Context represents each allocation site.

Typically offers quite precise context analysis

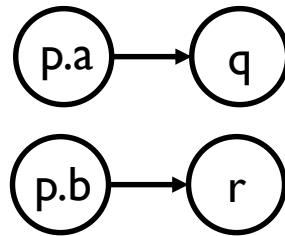
[Parameterized Object Sensitivity for Points-to and Side-Effect Analyses for Java;
Milanova et al. ISSTA 2002]

Field-Sensitivity of Analysis

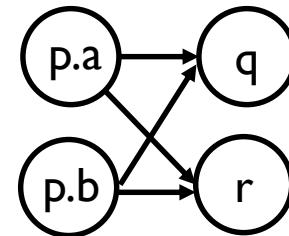
Def. A field-sensitive analysis is one that tracks distinct behavior for individual fields of a record type. Otherwise, it is field-insensitive

```
int f(T q, T r) {  
    p.a = &q;  
    p.b = &r;  
}
```

Field Sensitive



Field Insensitive



Challenges

- **Complexity:** For certain analysis techniques, converts linear representation to worse (perhaps even exponential)
- **Non-type-safe programs:** May have to track behavior at every byte offset within the structure (not just each field)

Flow Insensitive Algorithms

3 popular algorithms

- Any address
- Andersen, 1994
- Steensgard, 1996

Acceptable precision in practice for **compiler optimization**, however perhaps **insufficient for static analysis** approaches for security, reliability, or bug finding

Any Address Analysis

- ***Single points-to set***: contains all variables whose address is taken, passed by reference, etc.
- ***Any pointer may point to any variable*** in this set
- Simple, fast, linear-time algorithm
- Common choice for function pointers, and for global variables, e.g., for initial call graph
- Can refine with splitting by types

Example 1

```
void main() {  
    T *p, *q, *r;  
    T t;
```

```
o1:p = new T;
```

```
    q = &t;
```

```
    r = q;
```

```
}
```

```
// {p} -> {o1}
```

```
// {p,q} -> {o1, t}
```

```
// {p,q,r} -> {o1, t}
```

Andersen's Algorithm

- Generally the most precise flow- and context-insensitive algorithm
- Compute a single points-to graph for entire program
- Refinement by Burke: Separate points-to graph for each function
- Cost is $O(n^3)$ for program with n assignments
 - McAlister, On the complexity analysis of static analyses (SAS'99)
 - Sridharan and Fink, The Complexity of Andersen's Analysis in Practice (SAS'09)

Andersen's Algorithm: **Conceptual**

Initialize: Points-to graph with a separate node per variable

Iterate until convergence:

At each statement, evaluate the appropriate rule:

Form

$p = \&x$

$p = q$

$*p = q$

$p = *q$

Action

Add $p \rightarrow x$

$\forall x$: if $q \rightarrow x$, add $p \rightarrow x$

$\forall x, r$: if $q \rightarrow x$ and $p \rightarrow r$, add $r \rightarrow x$

$\forall x, r$: if $q \rightarrow x$ and $x \rightarrow r$, add $p \rightarrow r$

Andersen's Algorithm: **Actual**

1. Build initial "inclusion constraint graph" and initial points-to sets
2. Iterate until converged:
 - Update constraint graph for new points-to pairs
 - Update the points-to sets according to new constraints

Inclusion Constraint Graph: Add constraint for pointer assignments (pts is points-to set):

Name	Form	Constraint	Action
<i>Points-to pair</i>	$p = \&x$	$p \supseteq \{x\}$	$pts(p) \cup = \{x\}$
<i>Direct constraint</i>	$p = q$	$p \supseteq q$	$pts(p) \cup = pts(q)$
<i>Indirect constraint</i>	$*p = q$	$*p \supseteq q$	for $v \in pts(p)$: $pts(v) \cup = pts(q)$
<i>Indirect constraint</i>	$p = *q$	$p \supseteq *q$	for $v \in pts(q)$: $pts(p) \cup = pts(v)$

Example 1 Revisited

```
void main() {  
    T *p, *q, *r;  
    T t;
```

```
o1:p = new T;  
    q = &t;  
    r = q;  
}
```

```
// {p} -> {o1}  
// {p} -> {o1}, {q} -> {t}  
// {r} -> {t}
```

Form	Constraint	Action
$p = \&x$	$p \supseteq \{x\}$	$\text{pts}(p) \cup = \{x\}$
$p = q$	$p \supseteq q$	$\text{pts}(p) \cup = \text{pts}(q)$
$*p = q$	$*p \supseteq q$	for $v \in \text{pts}(p)$: $\text{pts}(v) \cup = \text{pts}(q)$
$p = *q$	$p \supseteq *q$	for $v \in \text{pts}(q)$: $\text{pts}(p) \cup = \text{pts}(v)$

Example 2

```
void f(int i) {  
    T *p, *q, *r;
```

```
    o1:p = new T;
```

```
    // {p} -> {o1}
```

```
    o2:q = new T;
```

```
    // {q} -> {o2}
```

```
        if (i>0)
```

```
            r = p;
```

```
            // {r} -> {o1}
```

```
        else
```

```
            r = q;
```

```
            // {r} -> {o2}
```

```
    }
```

Form	Constraint	Action
$p = \&x$	$p \supseteq \{x\}$	$\text{pts}(p) \cup = \{x\}$
$p = q$	$p \supseteq q$	$\text{pts}(p) \cup = \text{pts}(q)$
$*p = q$	$*p \supseteq q$	for $v \in \text{pts}(p)$: $\text{pts}(v) \cup = \text{pts}(q)$
$p = *q$	$p \supseteq *q$	for $v \in \text{pts}(q)$: $\text{pts}(p) \cup = \text{pts}(v)$

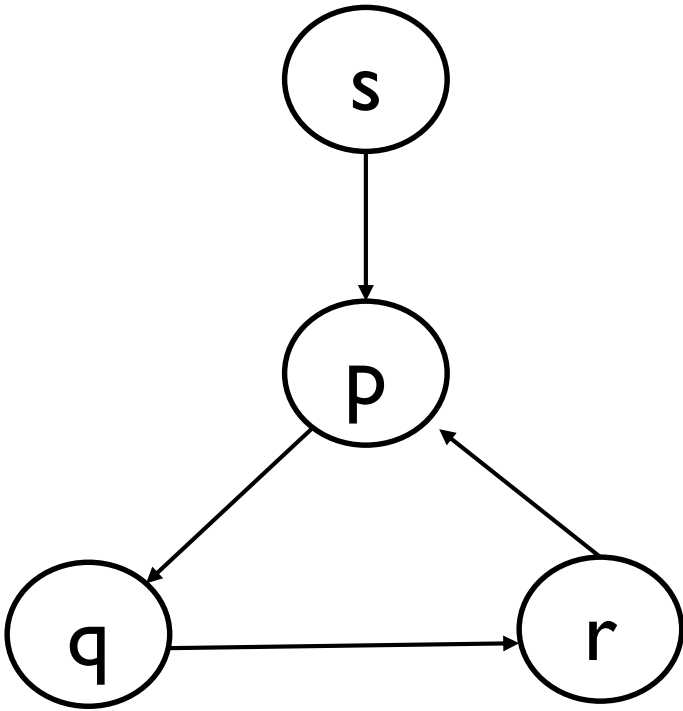
Example 3

<code>p = &a;</code>	<code>// p -> {a}</code>	<code>// p -> {a}</code>
<code>s = &p;</code>	<code>// s -> {p}</code>	<code>// s -> {p,q}</code>
<code>r = *s;</code>	<code>// r -> {a}</code>	<code>// r -> {a, b}</code>
<code>q = &b;</code>	<code>// q -> {b}</code>	<code>// q -> {t}</code>
<code>s = &q;</code>	<code>// s -> {p,q}</code>	<code>// s -> {p,q}</code>

Done?

Done?

Andersen's Algorithm: **Cycles**



Cycle in constraint graph:

$\text{pts}(p) \supseteq \text{pts}(q) \supseteq \text{pts}(r) \supseteq \text{pts}(p)$
 $\Rightarrow \text{pts}(p) = \text{pts}(q) = \text{pts}(r) = \text{pts}(p)$
 \Rightarrow No need to propagate points-to
pairs around such cycles!

Andersen's Algorithm: Cycles

Cycle in constraint graph:

$$\text{pts}(p) \supseteq \text{pts}(q) \supseteq \text{pts}(r) \supseteq \text{pts}(p)$$

$$\Rightarrow \text{pts}(p) = \text{pts}(q) = \text{pts}(r) = \text{pts}(p)$$

\Rightarrow No need to propagate points-to pairs around such cycles!

Offline cycle elimination:

- Find cycles due to direct pointer copies (direct constraints)
- Collapse each cycle into a single node, reduces size of constraint graph
- But many more cycles can be induced by indirect constraint edges: we need cycle elimination during transitive closure ("*online*")

"Off-line Variable Substitution for Scaling Points-To Analysis," Rountev and Chandra, PLDI'00.

Online cycle elimination:

- Fähndrich, Foster, Aiken and Su (PLDI '98): Cycle elimination is essential for scalability.
- Heintze and Tardieu (PLDI 2001): "A million lines of code per second."
- Hardekopf and Lin (PLDI 2007)

Steensgard's Algorithm

Unification:

- Conceptually: restrict every node to only one outgoing edge (on the fly)
- If $p \rightarrow x$ and $p \rightarrow y$, merge x and y (“unify”)
- All objects “pointed to” by p one equivalence class

$A = \&B$

$B = \&C$

$A = \&D$

$D = \&E$



Steensgard's Algorithm

Unification: Conceptually: restrict *every node to only one outgoing edge* (on the fly)

- If $p \rightarrow x$ and $p \rightarrow y$, merge x and y (“unify”)
- All objects “pointed to” by p form one equivalence class

Algorithm

1. For each statement, merge points-to sets:
 - $p = q$: Merge two equivalence classes (p 's and q 's targets)
Less expensive than computing points-to iterations
This may cause other nodes to collapse!
2. Use Tarjan's “union-find” (disjoint-set) data structure to record equivalence classes

Steensgard's Algorithm

“Union-find” aka Disjoint Set data structure:

- Splits the set of elements into disjoint partitions
- Maintains the partition with every addition
- Operations:
 - Find(x): follows parent pointers from x until reaching root (i.e. finds the set containing x)
 - Union(x,y): 1) finds the roots of x,y; 2) merges the trees by connecting the root nodes. (i.e. merges the sets)
- Properties: addition and merge of sets in near constant time, i.e. $\alpha(n)$ – inverse Ackerman func. $\alpha(n) < 4$ even for large n.

Consequence for Steensgard's analysis:

- Non-iterative algorithm, almost-linear running time: $O(n\alpha(n))$
- Like Andersen, single solution for the entire program

Steensgard vs. Andersen

Consider assignment $p = q$, i.e., only p is modified, not q .

Subset-based Algorithms (Anderson's algorithm is an example)

- Add a constraint: Targets of q must be subset of targets of p
- Graph of such constraints is also called “inclusion constraint graphs”
- Enforces unidirectional flow from q to p

Unification-based Algorithms (Steensgard is an example)

- Merge equivalence classes: targets of p and q must be identical
- Assumes bidirectional flow from q to p and vice-versa

In-between solutions:

- Unification-based Pointer Analysis with Directional Assignment, Das, PLDI 2000 – exploits the semantics of C; uses Andersen for top pointers, Steensgard elsewhere

Alias Analysis

- Alias analysis is a common client of pointer (points-to) analysis
 - **Pointer Analysis:** What objects does each pointer points to?
 - **Alias Analysis:** Can two pointers point to the same location? (i.e., it is possible that $*p = *q$)
- Once we have performed the pointer analysis, it is trivial to compute alias analysis (but not vice versa)
- **Two pointers p and q may alias if $\text{points-to}(p) \cap \text{points-to}(q) \neq \emptyset$**

Which Pointer Analysis To Use?

Hind & Pioli, ISSTA, Aug. 2000

Compared 5 algorithms (4 flow-insensitive, 1 flow-sensitive):

- Any address
- Steensgard
- Anderson
- Burke (like Anderson, but separate solution per procedure)
- Choi et al. (flow-sensitive)

Metrics

1. Precision: number of alias pairs
2. Precision of important optimizations: MOD/REF, REACH, LIVE, flow dependences, constant prop.
3. Efficiency: analysis time/memory, optimization time/memory

Benchmarks: 23 C programs, including some from SPEC benchmarks

Which Pointer Analysis To Use?

I. Precision: (Table 2)

- Steensgard much better than Any-Address (6x on average)
- Anderson/Burke significantly better than Steensgard (about 2x)
- Choi negligibly better than Anderson/Burke

Table 2: Mod and Ref at pointer dereferences and all CFG nodes. No assignments through a pointer occur in compiler. “AT” = Address Taken, “St” = Steensgard’s, “A/B”= Andersen/Burke et al., “Ch” = Choi et al.

Name	Ptr Mod				Ptr Ref				Mod				Ref			
	AT	St	A/B	Ch	AT	St	A/B	Ch	AT	St	A/B	Ch	AT	St	A/B	Ch
allroots	3	2.00	1.00	–	3	2.00	1.38	–	.88	.85	.83	–	1.77	1.58	1.52	–
052.alvinn	14	1.00	–	–	14	1.00	–	–	1.17	.51	–	–	2.24	1.00	–	–
01.qbsort	12	2.00	1.50	–	12	1.76	–	–	.95	.46	–	–	3.39	.94	–	–
06.matx	15	3.00	2.22	–	15	3.25	3.12	–	1.09	.39	.34	–	1.48	.89	.88	–
15.trie	10	1.12	–	1.00	10	1.00	–	–	1.02	.58	–	.52	4.42	.88	–	–
04.bisect	14	1.15	–	–	14	1.00	–	–	2.57	.58	–	–	3.92	1.57	–	–
fixoutput	9	1.80	–	–	9	2.00	–	–	.74	.37	–	–	.78	.56	–	–
17.bintr	7	1.00	–	–	7	1.00	–	–	.62	.30	–	–	2.07	.71	–	–
anagram	17	1.00	–	–	17	1.10	–	–	.90	.45	–	–	2.38	1.02	–	–
ks	17	1.90	1.86	1.62	17	1.79	–	1.74	1.70	.56	.55	.53	3.76	1.35	–	1.34
05.eks	12	1.22	–	–	12	1.02	–	–	1.83	.50	–	–	3.54	1.63	–	–
08.main	13	6.00	3.27	2.61	13	5.14	4.61	3.59	1.76	.63	–	–	5.35	1.32	1.30	1.25
09.vor	19	1.85	1.35	1.32	19	1.92	1.68	1.60	2.04	.63	.62	–	7.91	1.40	1.34	–
loader	47	3.77	2.23	–	47	2.09	1.36	–	5.08	.90	.73	–	9.72	1.78	1.30	–
129.compress	13	1.40	1.07	–	13	2.26	1.11	–	1.68	.80	.78	–	1.66	1.29	1.28	–
ft	10	2.87	1.80	1.72	10	2.66	2.53	2.39	2.14	.90	.74	.73	2.58	1.31	–	1.28
football	32	6.00	2.10	–	32	3.26	1.54	–	1.37	.70	.61	–	4.55	1.83	1.65	–
compiler	0	–	–	–	10	1.00	–	–	3.38	–	–	–	4.45	4.44	–	–
assembler	87	1.24	2.21	–	87	15.14	2.11	–	1.21	1.88	.87	–	15.07	4.09	1.47	–
yacr2	48	1.14	1.11	–	48	1.08	1.02	–	5.32	.53	.52	–	7.80	1.65	–	–
simulator	87	3.16	2.05	–	87	3.95	1.86	–	6.82	.62	.57	–	8.21	1.21	1.06	–
flex	56	5.37	1.78	–	56	5.09	2.03	2.01	5.97	1.60	1.18	–	1.55	3.89	3.44	3.43
099.go	154	42.68	13.64	–	154	51.39	17.03	–	7.31	5.87	3.94	–	4.47	3.98	3.13	–
Average	30.26	4.03	2.06	2.02	30.70	4.87	2.35	2.29	2.50	1.04	0.871	0.867	4.48	1.75	1.540	1.536

Which Pointer Analysis To Use?

2. MOD/REF precision: (Table 2)

- Steensgard much better than Any-Address (2.5x on average)
- Anderson/Burke significantly better than Steensgard (15%)
- Choi very slightly better than Anderson/Burke (1%)

Table 2: Mod and Ref at pointer dereferences and all CFG nodes. No assignments through a pointer occur in compiler. “AT” = Address Taken, “St” = Steensgard’s, “A/B”= Andersen/Burke et al., “Ch” = Choi et al.

Name	Ptr Mod				Ptr Ref				Mod				Ref			
	AT	St	A/B	Ch	AT	St	A/B	Ch	AT	St	A/B	Ch	AT	St	A/B	Ch
allroots	3	2.00	1.00	–	3	2.00	1.38	–	.88	.85	.83	–	1.77	1.58	1.52	–
052.alvinn	14	1.00	–	–	14	1.00	–	–	1.17	.51	–	–	2.24	1.00	–	–
01.qbsort	12	2.00	1.50	–	12	1.76	–	–	.95	.46	–	–	3.39	.94	–	–
06.matx	15	3.00	2.22	–	15	3.25	3.12	–	1.09	.39	.34	–	1.48	.89	.88	–
15.trie	10	1.12	–	1.00	10	1.00	–	–	1.02	.58	–	.52	4.42	.88	–	–
04.bisect	14	1.15	–	–	14	1.00	–	–	2.57	.58	–	–	3.92	1.57	–	–
fixoutput	9	1.80	–	–	9	2.00	–	–	.74	.37	–	–	.78	.56	–	–
17.bintr	7	1.00	–	–	7	1.00	–	–	.62	.30	–	–	2.07	.71	–	–
anagram	17	1.00	–	–	17	1.10	–	–	.90	.45	–	–	2.38	1.02	–	–
ks	17	1.90	1.86	1.62	17	1.79	–	1.74	1.70	.56	.55	.53	3.76	1.35	–	1.34
05.eks	12	1.22	–	–	12	1.02	–	–	1.83	.50	–	–	3.54	1.63	–	–
08.main	13	6.00	3.27	2.61	13	5.14	4.61	3.59	1.76	.63	–	–	5.35	1.32	1.30	1.25
09.vor	19	1.85	1.35	1.32	19	1.92	1.68	1.60	2.04	.63	.62	–	7.91	1.40	1.34	–
loader	47	3.77	2.23	–	47	2.09	1.36	–	5.08	.90	.73	–	9.72	1.78	1.30	–
129.compress	13	1.40	1.07	–	13	2.26	1.11	–	1.68	.80	.78	–	1.66	1.29	1.28	–
ft	10	2.87	1.80	1.72	10	2.66	2.53	2.39	2.14	.90	.74	.73	2.58	1.31	–	1.28
football	32	6.00	2.10	–	32	3.26	1.54	–	1.37	.70	.61	–	4.55	1.83	1.65	–
compiler	0	–	–	–	10	1.00	–	–	3.38	–	–	–	4.45	4.44	–	–
assembler	87	1.24	2.21	–	87	15.14	2.11	–	1.21	1.88	.87	–	15.07	4.09	1.47	–
yacr2	48	1.14	1.11	–	48	1.08	1.02	–	5.32	.53	.52	–	7.80	1.65	–	–
simulator	87	3.16	2.05	–	87	3.95	1.86	–	6.82	.62	.57	–	8.21	1.21	1.06	–
flex	56	5.37	1.78	–	56	5.09	2.03	2.01	5.97	1.60	1.18	–	1.55	3.89	3.44	3.43
099.go	154	42.68	13.64	–	154	51.39	17.03	–	7.31	5.87	3.94	–	4.47	3.98	3.13	–
Average	30.26	4.03	2.06	2.02	30.70	4.87	2.35	2.29	2.50	1.04	0.871	0.867	4.48	1.75	1.540	1.536

Which Pointer Analysis To Use?

3. Analysis cost: (Table 5)

- Any-Address, Steensgard extremely fast
- Anderson/Burke about 30x slower
- Choi about 2.5x slower than Anderson/Burke

Table 5: Analysis Time in Seconds

Name	Pointer Analysis					Clients					Total				
	AT	ST	An	Bu	Ch	AT	ST	An	Bu	Ch	AT	ST	An	Bu	Ch
allroots	0.01	0.01	0.02	0.03	0.02	0.08	0.08	0.12	0.10	0.09	0.09	0.09	0.14	0.13	0.11
052.alvinn	0.01	0.01	0.02	0.04	0.04	0.22	0.18	0.21	0.20	0.18	0.23	0.19	0.23	0.24	0.22
01.qbsort	0.01	0.01	0.05	0.13	0.30	0.11	0.06	0.09	0.09	0.07	0.12	0.07	0.14	0.22	0.37
06.matx	0.01	0.01	0.06	0.15	0.27	0.20	0.14	0.20	0.16	0.15	0.21	0.15	0.26	0.31	0.42
15.trie	0.01	0.01	0.11	0.19	0.16	0.07	0.07	0.12	0.07	0.08	0.08	0.08	0.23	0.26	0.24
04.bisect	0.01	0.01	0.03	0.05	0.05	0.15	0.10	0.13	0.11	0.10	0.16	0.11	0.16	0.16	0.15
fixoutput	0.01	0.01	0.01	0.06	0.06	0.07	0.06	0.10	0.08	0.07	0.08	0.07	0.11	0.14	0.13
17.bintr	0.01	0.01	0.06	0.05	0.09	0.06	0.07	0.06	0.07	0.05	0.07	0.08	0.12	0.12	0.14
anagram	0.01	0.01	0.09	0.21	0.17	0.34	0.30	0.34	0.30	0.30	0.35	0.31	0.43	0.51	0.47
ks	0.01	0.01	0.08	0.37	0.51	0.34	0.22	0.27	0.21	0.22	0.35	0.23	0.35	0.58	0.73
05.eks	0.01	0.01	0.08	0.17	0.26	0.70	0.65	0.74	0.66	0.68	0.71	0.66	0.82	0.83	0.94
08.main	0.01	0.01	0.66	1.12	1.44	1.14	0.92	1.03	0.95	0.94	1.15	0.93	1.69	2.07	2.38
09.vor	0.01	0.01	3.21	4.05	6.24	1.44	1.09	1.69	1.11	1.04	1.45	1.10	4.90	5.16	7.28
loader	0.01	0.01	0.62	0.59	1.88	2.18	1.84	2.18	1.47	1.46	2.19	1.85	2.80	2.06	3.34
129.compress	0.03	0.04	0.02	0.07	0.08	1.99	1.82	1.97	1.79	1.83	2.02	1.86	1.99	1.86	1.91
ft	0.01	0.01	0.53	1.44	2.27	0.43	0.23	0.40	0.27	0.25	0.44	0.24	0.93	1.71	2.52
football	0.01	0.01	1.34	0.89	1.40	10.38	10.28	11.65	8.50	8.45	10.39	10.29	12.99	9.39	9.85
compiler	0.01	0.01	0.08	0.67	0.64	1.94	2.06	3.09	2.02	2.16	1.95	2.07	3.17	2.69	2.80
assembler	0.01	0.02	5.68	1.77	4.81	6.17	5.21	6.92	3.52	3.26	6.18	5.23	12.60	5.29	8.07
yacr2	0.01	0.01	1.10	2.37	4.92	8.14	6.53	9.61	5.79	5.05	8.15	6.54	10.71	8.16	9.97
simulator	0.02	0.01	1.87	2.19	6.94	16.13	12.80	15.14	8.82	8.42	16.15	12.81	17.01	11.01	15.36
flex	0.02	0.02	4.20	11.59	36.88	44.31	30.84	33.18	29.05	28.26	44.33	30.86	37.38	40.64	65.14
099.go	0.73	0.62	9.38	4.39	9.27	98.96	83.82	74.42	73.18	72.54	99.69	84.44	83.80	77.57	81.81
Ratio to AT	1.00	0.90	29.60	32.92	79.49	1.00	0.81	0.84	0.71	0.69	1.00	0.82	0.98	0.87	1.09

Which Pointer Analysis To Use?

4. Total cost of analysis + optimizations: (Table 5)

- *the client analyses improved in efficiency as the pointer information more precise*
- Steensgard, Burke are 15% faster than Any-Address!
- Anderson is as fast as Any-Address!
- Choi only about 9% slower

Table 5: Analysis Time in Seconds

Name	Pointer Analysis					Clients					Total				
	AT	ST	An	Bu	Ch	AT	ST	An	Bu	Ch	AT	ST	An	Bu	Ch
allroots	0.01	0.01	0.02	0.03	0.02	0.08	0.08	0.12	0.10	0.09	0.09	0.09	0.14	0.13	0.11
052.alvinn	0.01	0.01	0.02	0.04	0.04	0.22	0.18	0.21	0.20	0.18	0.23	0.19	0.23	0.24	0.22
01.qbsort	0.01	0.01	0.05	0.13	0.30	0.11	0.06	0.09	0.09	0.07	0.12	0.07	0.14	0.22	0.37
06.matx	0.01	0.01	0.06	0.15	0.27	0.20	0.14	0.20	0.16	0.15	0.21	0.15	0.26	0.31	0.42
15.trie	0.01	0.01	0.11	0.19	0.16	0.07	0.07	0.12	0.07	0.08	0.08	0.08	0.23	0.26	0.24
04.bisect	0.01	0.01	0.03	0.05	0.05	0.15	0.10	0.13	0.11	0.10	0.16	0.11	0.16	0.16	0.15
fixoutput	0.01	0.01	0.01	0.06	0.06	0.07	0.06	0.10	0.08	0.07	0.08	0.07	0.11	0.14	0.13
17.bintr	0.01	0.01	0.06	0.05	0.09	0.06	0.07	0.06	0.07	0.05	0.07	0.08	0.12	0.12	0.14
anagram	0.01	0.01	0.09	0.21	0.17	0.34	0.30	0.34	0.30	0.30	0.35	0.31	0.43	0.51	0.47
ks	0.01	0.01	0.08	0.37	0.51	0.34	0.22	0.27	0.21	0.22	0.35	0.23	0.35	0.58	0.73
05.eks	0.01	0.01	0.08	0.17	0.26	0.70	0.65	0.74	0.66	0.68	0.71	0.66	0.82	0.83	0.94
08.main	0.01	0.01	0.66	1.12	1.44	1.14	0.92	1.03	0.95	0.94	1.15	0.93	1.69	2.07	2.38
09.vor	0.01	0.01	3.21	4.05	6.24	1.44	1.09	1.69	1.11	1.04	1.45	1.10	4.90	5.16	7.28
loader	0.01	0.01	0.62	0.59	1.88	2.18	1.84	2.18	1.47	1.46	2.19	1.85	2.80	2.06	3.34
129.compress	0.03	0.04	0.02	0.07	0.08	1.99	1.82	1.97	1.79	1.83	2.02	1.86	1.99	1.86	1.91
ft	0.01	0.01	0.53	1.44	2.27	0.43	0.23	0.40	0.27	0.25	0.44	0.24	0.93	1.71	2.52
football	0.01	0.01	1.34	0.89	1.40	10.38	10.28	11.65	8.50	8.45	10.39	10.29	12.99	9.39	9.85
compiler	0.01	0.01	0.08	0.67	0.64	1.94	2.06	3.09	2.02	2.16	1.95	2.07	3.17	2.69	2.80
assembler	0.01	0.02	5.68	1.77	4.81	6.17	5.21	6.92	3.52	3.26	6.18	5.23	12.60	5.29	8.07
yacr2	0.01	0.01	1.10	2.37	4.92	8.14	6.53	9.61	5.79	5.05	8.15	6.54	10.71	8.16	9.97
simulator	0.02	0.01	1.87	2.19	6.94	16.13	12.80	15.14	8.82	8.42	16.15	12.81	17.01	11.01	15.36
flex	0.02	0.02	4.20	11.59	36.88	44.31	30.84	33.18	29.05	28.26	44.33	30.86	37.38	40.64	65.14
099.go	0.73	0.62	9.38	4.39	9.27	98.96	83.82	74.42	73.18	72.54	99.69	84.44	83.80	77.57	81.81
Ratio to AT	1.00	0.90	29.60	32.92	79.49	1.00	0.81	0.84	0.71	0.69	1.00	0.82	0.98	0.87	1.09

Analysis Scalability

	Equality-based	Subset-based	Flow-sensitive
Context-insensitive	<ul style="list-style-type: none"> • Wehl [32] 1980: < 1 KLOC first paper on pointer analysis • Steensgaard [31] 1996: 1+ MLOC first scalable pointer analysis 	<ul style="list-style-type: none"> • Andersen [1] 1994: 5 KLOC • Fähndrich et al. [7] 1998: 60 KLOC • Heintze and Tardieu [11] 2001: 1 MLOC • Berndt et al. [2] 2003: 500 KLOC first to use BDDs 	<ul style="list-style-type: none"> • Choi et al. [5] 1993: 30 KLOC
Context-sensitive	<ul style="list-style-type: none"> • Fähndrich et al. [8] 2000: 200K 	<ul style="list-style-type: none"> • Whaley and Lam [35] 2004: 600 KLOC cloning-based BDDs 	<ul style="list-style-type: none"> • Landi and Ryder [19] 1992: 3 KLOC • Wilson and Lam [37] 1995: 30 KLOC • Whaley and Rinard [36] 1999: 80 KLOC

Derek Rayside, Points-To Analysis (Summary), 2005

<https://www.cs.utexas.edu/~pingali/CS395T/2012sp/lectures/points-to.pdf>

More recent: Flow-Sensitive Pointer Analysis for Millions of Lines of Code
Hardekopf and Lin (CGO'11)

Advanced Techniques

- **Shape Analysis:** discovers and reasons about dynamically allocated data structures (e.g., lists, trees, heaps)
- **Escape Analysis:** computes which program locations can access a pointer (across function boundaries)
- **Datalog:** Declarative, constraint-based approach to specify analysis, offers pretty good scalability

Pointer Analysis; Yannis Smaragdakis; George Balatsouras, Now Publishing, 2015

Datalog

Datalog: declarative language with Prolog-like notation

Elements: **atoms** of the form $p(X_1, X_2, \dots, X_n)$

- p is a predicate
- X_1, X_2, \dots, X_n are variables or constants

Ground atoms: predicate with only constant arguments

- Its value is either true or false

Rules: $H :- B_1 \& B_2 \& \dots \& B_n$

- H is an **atom**, $B_1 \dots B_n$ are **atoms** or **negations** of atoms
- $:-$ is “if” --- so H is valid if all $B_1 \dots B_n$ are valid

Datalog program is a collection of rules. The program is applied to a set of ground atoms. The result is the set of ground atoms inferred by applying the rules until fixpoint

Datalog Example

Simple Datalog program (from Dragon book):

```
path(X,Y) :- edge (X,Y)
```

```
path(X,Y) :- path (X,Z) & path (Z,Y)
```

The meaning of the program: A single edge is a path; a path also exist if there is a path between the start point and some other point, and that other point and the end point.

Consider this example:

- True ground atoms: edge(1,2), edge(2,3), edge(3,4)
- Infer path(1,2), path(2,3), path(3,4) using rule #1
- Infer composite paths using successive application of rule #2

Flow-Insensitive Pointer Analysis

(Dragonbook) Compute:

- **Pts**(V, H) – the variable V can point to heap object H
- **Hpts**(H, F, G) – field F of heap object H points to heap object G

Rules constructed by traversing the program:

1. **Pts(V, H) :- “H: V = malloc”**

V points to heap loc H if it is allocated at H (say we use line number calling)

2. **Pts(V, H) :- “V = W” & Pts (W, H)**

V points to H if V points to W and W points to H

3. **Hpts(H, F, G) :- “V.F = W” & Pts(W, G) & Pts(V, H)**

In stmt V.F=W, field F of object H points to object G if ptr W points to G and ptr V points to H

4. **Pts(V, H) :- “V = W.F” & Pts(W, G) & Hpts(G, F, H)**

In stmt V=W.F, V points to H if W points to G and field F of G points to H

Context-Sensitive Pointer Analysis

First compute:

- **Pts**(V, C, H) – the variable V in context C can point to heap object H
- **Hpts**(H, F, G) – field F of heap object H points to heap object G
- **CSinvokes**(S, C, M, D) – the calls site S in context C calls the D context of M

Rules constructed by traversing the program:

1. **Pts(V, C, H) :- “H: V = malloc” & CSinvokes(H, C, _, _)**
2. **Pts(V, C, H) :- “V = W” & Pts(W, C, H)**
3. **Hpts(H, F, G) :- “V.F = W” & Pts(W, C, G) & Pts(V, C, H)**
4. **Pts(V, C, H) :- “V = W.F” & Pts(W, G) & Hpts(G, F, H)**
5. **Pts(V, D, H) :- CSinvokes(H, C, M, D) & formal(M, D, V) & actual(S, C, W) & pts(W, C, H)**

If the call site S in context C calls method M of context D, then the formal parameters in method M of context D can point to the objects pointed to by the actual params in C