

*Precise Complexity Guarantees for Pointer Analysis via Datalog with Extensions**

K. TUNCAY TEKLE, YANHONG A. LIU

Computer Science Department, Stony Brook University, Stony Brook, NY, USA

(e-mail: tuncay,liu@cs.stonybrook.edu)

Abstract

Pointer analysis is a fundamental static program analysis for computing the set of objects that an expression can refer to. Decades of research has gone into developing methods of varying precision and efficiency for pointer analysis for programs that use different language features, but determining precisely how efficient a particular method is has been a challenge in itself.

For programs that use different language features, we consider methods for pointer analysis using Datalog and extensions to Datalog. When the rules are in Datalog, we present the calculation of precise time complexities from the rules using a new algorithm for decomposing rules for obtaining the best complexities. When extensions such as function symbols and universal quantification are used, we describe algorithms for efficiently implementing the extensions and the complexities of the algorithms.

KEYWORDS: Datalog, function symbols, universal quantification, computational complexity, static program analysis, pointer analysis, alias analysis

1 Introduction

Pointer analysis is a static program analysis for computing the set of objects that an expression can refer to. It is a fundamental analysis used for many applications, e.g., debugging (Shapiro and Horwitz 1997), performance analysis (Ghiya et al. 2001), dataflow analysis (Shapiro and Horwitz 1997), parallelism (Wilson and Lam 1995; Pearce et al. 2007), common subexpression elimination (Diwan et al. 1998; Ghiya and Hendren 1998), optimization by incrementalization (Gorbovitski et al. 2010), and detection of security vulnerabilities (Avots et al. 2005). Consider the following program fragment in an object-oriented programming language:

```
void foo() {  
    Object o1 = new Object();  
    Object o2;  
    if (...) o2 = id(o1);  
    else o2 = new Object();  
}  
void id(Object o) { return o; }
```

For each variable $o1$, $o2$, and o , a pointer analysis may aim to find the set of objects the variable *may* point to. We call each such set a *may-point-to* set. A may-point-to analysis is *sound* if each may-point-to set contains all objects that the variable may point to at runtime. A sound

* This work was supported in part by NSF under grants CCF-1414078, IIS-1447549, CCF-1248184, CCF-0964196, and ONR under grant N000141512208.

may-point-to analysis is *precise* if no may-point-to set contains more objects than the set of objects that the variable may point to at runtime. Another form of analysis is *must-point-to* analysis, which finds the set of objects that each variable *must* point to. Precise pointer analysis is undecidable (Landi 1992; Ramalingam 1994). Therefore, pointer analysis methods approximate the results, providing different tradeoffs between the precision of the results and the efficiency of the method (Hind and Pioli 2000) while preserving soundness. Whereas may-point-to analysis is an overapproximation, must-point-to analysis is an underapproximation.

From the simple program fragment above, it can be seen that the precision-efficiency tradeoff involves the consideration of control flows, procedures, calling contexts, objects, and other language features such as arrays; leading to a variety of analyses. The worst-case time complexities of existing analyses vary from almost linear (Steensgaard 1996) to doubly exponential (Sagiv et al. 1998). However, such worst-case complexities are often not a true indication of analysis time; many researchers provide empirical performance results for their algorithms, and many papers have been written on which pointer analysis one should use, e.g., (Hind and Pioli 2000), including one questioning whether we have solved the pointer analysis problem yet (Hind 2001).

A recent survey (Smaragdakis and Balatsouras 2015) presents existing work on logical specifications of pointer analysis methods in the declarative language *Datalog* and its extensions. Datalog specifications allow expressing the precision aspects of pointer analyses concisely, while ensuring that the analyses are performed in polynomial time, because evaluation of rules in Datalog is guaranteed to be polynomial time. However, just as the worst-case complexities of existing analyses are not a true indication of analysis time, worst-case polynomial time for evaluation of Datalog is not sufficient for understanding actual running times.

In this paper, we consider all the different analyses presented in the survey (Smaragdakis and Balatsouras 2015) expressed in Datalog and its extensions, and study the time complexity of each analysis by using and extending a systematic method for calculating the time complexities for optimal bottom-up evaluation of Datalog rules (Liu and Stoller 2009). To obtain the best complexity, we give a new algorithm for rule decomposition. All analyses require handling rules with many hypotheses, and some require handling extension to Datalog with negation, function symbols, and universal quantification. In each case, we describe the method for handling the extensions and calculating the time complexities. Our methods can be readily used for analyzing other program analyses expressible in Datalog and similar extensions.

2 Language and preliminaries

In this section, we describe Datalog, an optimal method for evaluating a set of Datalog rules with at most two hypotheses each, and a method for calculating the time complexity of the evaluation.

Datalog. *Datalog* is a language for defining rules, facts, and queries, where rules can be used with facts to answer queries. A Datalog rule is of the form:

$$p(a_1, \dots, a_k) \leftarrow p_1(a_{11}, \dots, a_{1k_1}), \dots, p_h(a_{h1}, \dots, a_{hk_h}).$$

where h is a finite natural number, each p_i (respectively p) is a predicate of finite number k_i (respectively k) arguments, each a_{ij} and a_i is either a constant or a variable, and each variable in the arguments of p must also be in the arguments of some p_i . If $h = 0$, then each a_i must be a constant, in which case $p(a_1, \dots, a_k)$ is called a *fact*. For the rest of the paper, “rule” refers only to the case where $h \geq 1$, in which case each $p_i(a_{i1}, \dots, a_{ik_i})$ is called a *hypothesis*, and $p(a_1, \dots, a_k)$ is

called the *conclusion*. For rules with the same hypotheses but different conclusions, we use the shorthand of writing one rule with the same hypotheses but with comma-separated conclusions.

The meaning of a set of rules and facts is the set of facts that are given or can be inferred using the rules.

Terminology. An *IDB (intensional database) predicate* is a predicate defined by rules, and an *EDB (extensional database) predicate* is a predicate for which no rules exist, and only facts are given. An *IDB (EDB) hypothesis* is a hypothesis whose predicate is an IDB (EDB) predicate.

For complexity calculation, we use the following notations.

- $\#p$: number of facts of predicate p , called *size of p* .
- $\#\text{dom}(p, i)$: size of the domain from which the i th argument of predicate p takes value.
- $\#p.i$: number of values actually taken by the i th argument of the facts of predicate p (given or inferred).
- $\#p.i_1, \dots, i_n / j_1, \dots, j_m$: maximum number of combinations of different values actually taken by the i_1, \dots, i_n th arguments of the facts of predicate p (given or inferred), given any fixed value for the j_1, \dots, j_m th arguments.

We assume that hash tables, tries or similar data structures are used so that operations involving a single element of a set take $O(1)$ time.

2.1 Bottom-up evaluation and complexity calculation

Bottom-up evaluation starts with given facts, infers new facts from conclusions of rules whose hypotheses match existing facts, and does so repeatedly until all facts are inferred. In this paper, we use the bottom-up evaluation method of (Liu and Stoller 2009). The time complexity incurred by each rule using this method is bound by the *number of firings* of the rule—the number of combinations of facts that make all hypotheses true. The best complexity is the minimum among all possible decompositions of the rule into rules with at most two hypotheses. However, the number of decompositions of a rule is worse than exponential in the number of hypotheses.

In this subsection, we summarize how to compute the optimal time complexity incurred by a rule with at most two hypotheses with the method in (Liu and Stoller 2009). Note that the number of rules and the arities of predicates are considered constants, not affecting the asymptotic analysis. In the next section, we describe a heuristic algorithm for decomposing rules so that each rule has at most two hypotheses and then calculating the optimal complexity for the decomposition.

There are two forms of rules when the rules are limited to two hypotheses. When a rule has one hypothesis, it is of the form: $p(\dots) \leftarrow q(\dots)$. The number of times this rule can fire is the number of facts of q , therefore the time complexity incurred by this rule is $O(\#q)$. In fact, we can omit the complexity of such rules, because (i) if q is an EDB predicate, then all of its facts need to be read in, therefore $O(\#q)$ cost is already incurred by the reading of the input, (ii) if q is an IDB predicate, then its size would be bound by the complexity of the rules that infer its facts, and therefore that complexity would already have been included by the complexity of the rules inferring its facts.

When a rule has two hypotheses, it is of the form: $p(\dots) \leftarrow q(x_1, \dots, x_n), r(y_1, \dots, y_m)$. To calculate the number of firings, we can first think of processing the facts of q and matching them with facts of the second hypothesis such that the common variables in the hypotheses take the same value. Therefore, only the variables of the second hypothesis not in the first can take

values for each fact of q . We use C_{12} to denote the set of integers j in $[1, \dots, m]$ such that y_j is a variable common to both hypotheses, then the complexity is bounded by $O(\#q \times \#r \cdot i \notin C_{12} / j \in C_{12})$. Analogously, we can think of processing the facts of r and matching them with facts of the first hypothesis. Defining C_{21} analogously to be the set of integers j in $[1, \dots, n]$ such that x_j is a common variable, the complexity is also bounded by $O(\#r \times \#q \cdot i \notin C_{21} / j \in C_{21})$. Since both bounds are upper bounds on the number of firings, then the complexity is bounded by the minimum of the two: $O(\min(\#q \times \#r \cdot i \notin C_{12} / j \in C_{12}, \#r \times \#q \cdot i \notin C_{21} / j \in C_{21}))$. For example, for the rule $p(x, z) \leftarrow q(x, y), r(y, z)$, the complexity is $O(\min(\#q \times \#r \cdot 2/1, \#r \times \#q \cdot 1/2))$.

3 Handling many hypotheses—applied to pointer analyses for different language features

In general, rules may have many hypotheses. In this section, we describe an algorithm for decomposing rules so that the resulting set of rules has at most two hypotheses and achieves the best complexity among all possible decompositions. We then apply our algorithm coupled with the complexity calculation of bottom-up evaluation to calculate the complexity of may-point-to analysis for object-oriented languages. The algorithm applies also to analyses of other advanced language features, including procedures, arrays, and exceptions as shown in Appendix A.

3.1 An algorithm for decomposing rules with many hypotheses

Given a set of rules where some rules have more than two hypotheses, each such rule can be *decomposed* so that a set of rules with the same meaning is produced where each rule has two hypotheses. To decompose a rule R , we (1) select two hypotheses h_1 and h_2 of R ; (2) create a new intermediate rule R' whose hypotheses are h_1 and h_2 , and whose conclusion c is a new, intermediate predicate whose arguments are the variables occurring in h_1 or h_2 that are also elsewhere in R ; (3) replace h_1 and h_2 in R with c ; and (4) repeat steps (1)–(3) until R has only two hypotheses.

How two hypotheses are selected at step (1) of each iteration so that the running time of the resulting set of rules is minimized is analogous to the join-order optimization problem on relational database queries, which is well studied with many heuristic algorithms (Selinger et al. 1979; Steinbrunn et al. 1997). However, most heuristics assume that the sizes of predicates are known in advance, since they consider sizes of only EDB predicates. We propose a new heuristic algorithm that is deterministic and well-suited to Datalog applications. Our algorithm is presented below, where for each substep, we give the rationale.

- If there is any pair of hypotheses such that the variables of one hypothesis is a subset of the variables of the other, we select any one such pair. This ensures that the intermediate rule has no added asymptotic complexity.
- Otherwise, apply the following steps in order, each step applied to the set of pairs selected so far, starting with the set of all pairs, until a unique pair is selected:
 - (a.i) In a rule, a variable is called *removable* for two hypotheses if the variable only appears in those two hypotheses and nowhere else. We select all pairs of hypotheses with the maximum number of removable variables.
 - (a.ii) If domain size information is available, we multiply the sizes of the domains of each removable variable for each pair of hypotheses, and select all pairs with the maximum product. Steps (a.i) and (a.ii) help to minimize the matching of the different values of the removable variables with other hypotheses in the rest of the rule.

- (b) We select all pairs of hypotheses that contain the maximum number of shared variables between the hypotheses in the pair. This helps to minimize the number of facts iterated over during the evaluation of the intermediate rule to be created.
- (c.i) We select all pairs of hypotheses that contain the maximum number of EDB hypotheses. This helps to best understand the complexity of the intermediate rule because sizes of EDB predicates are input parameters.
- (c.ii) We select all pairs of hypotheses in which the product of the sizes of the EDB hypotheses in the pair is the minimum. This helps to minimize the cost of the intermediate rule.
- (d) As a last resort, we select the pair of the leftmost two hypotheses.

Next, we show applications of this algorithm coupled with evaluation and complexity calculation.

3.2 Andersen’s pointer analysis for object-oriented languages and its complexity

Pointer analysis comes in many flavors depending on what it takes into account. An *intraprocedural* analysis only considers a single procedure. An *interprocedural* analysis considers multiple procedures and interactions among them. A *flow-insensitive* analysis does not take control flows into account, whereas a *flow-sensitive* analysis does and produces a may-point-to or must-point-to set for each variable at each program point. A *context-sensitive* analysis takes calling contexts into account, and produces a may-point-to or must-point-to set for each variable for each possible calling context. A *context-insensitive* analysis does not.

For large programs, it is generally understood that flow-sensitivity and context-sensitivity are not feasible. The most well-known flow- and context-insensitive pointer analysis was developed by Andersen (Andersen 1994), and it is considered to offer a sweet spot between precision and efficiency (Hardekopf and Lin 2007). Andersen formulates a may-point-to analysis in terms of type theory, and the formulation corresponds directly to a logical specification in Datalog. This subsection considers an intraprocedural pointer analysis in Datalog for an object-oriented language based on Andersen’s analysis as described in (Smaragdakis and Balatsouras 2015). We show the decomposition of the rules and calculate the precise time complexity.

Predicates and rules for Andersen’s analysis for OO languages. Each statement has a corresponding fact, shown below. For the first statement, h , called a *heap abstraction*, is a new constant created as an abstraction for the set of possible heap objects created by `new` when executing the statement, and m , not used in this intraprocedural analysis but used in later analyses, is the method containing the statement:

<code>v = new Obj()</code>	<code>alloc(v,h,m)</code>
<code>v = v2</code>	<code>move(v,v2)</code>
<code>v.f = v2</code>	<code>store(v,f,v2)</code>
<code>v = v2.f</code>	<code>load(v,v2,f)</code>

For example, method `foo` in Section 1 has three facts: `alloc(o1,h1,foo)`, `alloc(o2,h2,foo)`, and `move(o2,o1)`, where $h1$ and $h2$ are fresh constants.

The analysis defines the following two predicates and infers facts of them using the Datalog rules (R1)–(R4) in Fig. 1; additional explanations can be found in (Smaragdakis and Balatsouras 2015).

```

v_pt(v,h) ← alloc(v,h,m). (R1)
v_pt(v,h) ← move(v,v2), v_pt(v2,h). (R2)
v_pt(v,h) ← load(v,v2,f), v_pt(v2,h2), f_pt(h2,f,h). (R3)
f_pt(h,f,h2) ← store(v,f,v2), v_pt(v,h), v_pt(v2,h2). (R4)

int1(v,f,h2) ← load(v,v2,f), v_pt(v2,h2). (R3/1) // two decomposed rules for (R3)
v_pt(v,h) ← int1(v,f,h2), f_pt(h2,f,h). (R3/2) //
int2(f,v2,h) ← store(v,f,v2), v_pt(v,h). (R4/1) // two decomposed rules for (R4)
f_pt(h,f,h2) ← int2(f,v2,h), v_pt(v2,h2). (R4/2) //

```

Fig. 1. An intraprocedural may-point-to analysis in Datalog, (R1)–(R4), and decomposed rules

Table 1. Complexities for an intraprocedural may-point-to analysis

(R1)	$O(\#alloc)$
(R2)	$O(\min(\#move \times \#v_pt.2/1, \#v_pt \times \#move.1/2))$
(R3/1)	$O(\min(\#load \times \#v_pt.2/1, \#v_pt \times \#load.1,3/2))$
(R3/2)	$O(\min(\#int1 \times \#f_pt.3/1,2, \#f_pt \times \#int1.1/2,3))$
(R4/1)	$O(\min(\#store \times \#v_pt.2/1, \#v_pt \times \#store.2,3/1))$
(R4/2)	$O(\min(\#int2 \times \#v_pt.2/1, \#v_pt \times \#int2.1,3/2))$

- $v_pt(v,h)$: variable v may point to heap abstraction h
- $f_pt(h1,f,h2)$: heap abstraction $h1$ may have its field f pointing to heap abstraction $h2$

Decomposition of rules and complexity analysis. Using our algorithm for rule decomposition in Section 3.1, we decompose rules (R3) and (R4) as shown in Fig. 1. For our rule decompositions, we note that in the input programs, there are more program points than variables, and more variables than heap abstractions. For (R3), by algorithm step (a.ii), we select the first two hypotheses because $v2$ is the removable variable whose domain size is maximum. For (R4), we can select the first two hypotheses or the first and the third because v and $v2$ are removable and their domains are the same, so we select the leftmost two hypotheses that remove v .

We calculate precise complexities for the decomposed rules, and show the results in Table 1. For the rest of the paper, p is the number of program points, v is the number of variables, h is the number of heap abstractions (i.e., $\#alloc$), and f is the number of fields.

The sizes of IDB predicates are bounded by both the complexities of the rules inferring their facts, and the product of the sizes of the domains of their arguments. Sizes $\#int1$ and $\#int2$ are bounded by the complexities of (R3/1) and (R4/1) respectively, because each firing produces at most one new fact; they are also bounded by $O(v \times h \times f)$ based on the domains of their arguments. In all cases, $\#v_pt$ is bounded by $O(v \times h)$, and $\#f_pt$ is bounded by $O(h^2 \times f)$. These complexities can be factored in when calculating the overall complexities.

Now, we give some insight into the calculated complexities. If we consider the worst case when all predicates are maximized (i.e., they have facts for all possible combinations of their argument values), the complexity of this analysis would be $O(p \times h^2)$. Under various conditions, we can obtain better complexities. For example, if all variables point to a constant number of heap abstractions, i.e., $O(\#v_pt.2/1)=O(1)$, then the complexity would be $O(p \times h)$. If, in addition, all fields of variables point to a constant number of heap abstractions, i.e., $O(\#f_pt.3/1,2)=O(1)$, then the complexity would be linear in the program size, $O(p)$.

4 Handling rules with function symbols—applied to context-sensitive may-point-to analyses

A context-sensitive may-point-to analysis separates may-point-to sets for executions that map to different contexts, thereby increasing precision. One can consider different types of contexts such as call sites (Shivers 1991; Sharir and Pnueli 1981), objects (Milanova et al. 2005), and types (Smaragdakis et al. 2011). A rule-based model of context-sensitive analysis is presented in (Smaragdakis and Balatsouras 2015), but the rules contain function symbols, invalidating the polynomial-time evaluation guarantee for pure Datalog; different restrictions to the function symbols are provided to ensure polynomial-time evaluation. In this section, we consider the restriction for the most sophisticated analysis, show how to extend our evaluation to handle function symbols, and calculate the complexities.

Evaluation of Datalog with bounded-size terms. We extend Datalog so that arguments of predicates may be terms, where a *term* is either a constant, a variable, or a function symbol with arguments that are terms. We denote function symbols with uppercase letters¹, and require that each function symbol F be of fixed arity, so all occurrences of F take the same number of arguments. If a term is a function symbol F with arguments, we call it a *term of F* . The number of constants and function symbols in a term is called its *size*. The introduction of function symbols to Datalog rules makes the language Turing-complete, therefore invalidating complexity and termination guarantees (Schreye and Decorte 1994). We introduce a sufficient condition for detection of termination in the presence of function symbols, and discuss the evaluation and complexity when termination is guaranteed.

We say that a rule is *size-bounding* for F if the sizes of the terms of F in the conclusion are guaranteed to be no larger than the size of the term of F with the maximum size in the hypotheses. If a set of Datalog rules extended with function symbols is size-bounding for every function symbol, then bottom-up evaluation is guaranteed to terminate. Note that rules with no function symbols in the conclusion are size-bounding by definition.

Given a set of size-bounding Datalog rules, we perform bottom-up evaluation and calculate its complexity exactly as described before. However, the sizes of predicates and domains of predicate arguments need to be made more precise for calculating the number of firings, because they can take on terms as values. For a function symbol F , we define $\text{count}(F)$ to be the number of different terms of F that can appear during evaluation. If a size-bounding rule r has a term of F in the conclusion distinct from terms of F that appear in its hypotheses, then the contribution of r to $\text{count}(F)$ is bounded by the product of the domains of the variables that appear in the term of F in the conclusion of r . Therefore, $\text{count}(F)$ is bounded by the sum of such contributions in every rule. The size of the domain of the i th argument of a predicate p ($\#\text{dom}(p.i)$) is bounded by $\text{count}(F)$, if there is a rule whose conclusion's predicate is p and (i) its i th argument is a term of F , or (ii) there is a hypothesis of predicate q whose j th argument a_j is bounded by $\text{count}(F)$ and the i th argument of the conclusion is a_j .

2-call-site sensitive analysis with a 1-call-site sensitive heap. A pointer analysis is said to be n -call-site sensitive if it tracks the last n method calls leading to the execution of a statement, with an m -call-site sensitive heap if it tracks the last m method calls leading to the creation of a

¹ In logic programming, the converse is true; we use this notation to emphasize their presence.

heap object. Out of three context-sensitive analyses in (Smaragdakis and Balatsouras 2015), we consider the most complex one, a 2-call-site sensitive analysis with a 1-call-site sensitive heap. The following additional kinds of facts are used:

<code>vcall(v, s, p, m)</code>	virtual call <code>v.s(...)</code> is at program point <code>p</code> in method <code>m</code>
<code>htype(h, t)</code>	heap abstraction <code>h</code> has type <code>t</code>
<code>lookup(t, s, m)</code>	method with signature <code>s</code> of type <code>t</code> is <code>m</code>
<code>this(m, t)</code>	<code>this</code> variable for method <code>m</code> is <code>t</code>
<code>farg(m, n, a)</code>	method <code>m</code> 's <code>n</code> th formal argument is <code>a</code>
<code>aarg(p, n, a)</code>	program point <code>p</code> is a call whose <code>n</code> th actual argument is <code>a</code>
<code>fret(m, v)</code>	method <code>m</code> 's formal return variable is <code>v</code>
<code>aret(p, v)</code>	program point <code>p</code> is a call that assigns to actual return variable <code>v</code>
<code>astore(v1, v2)</code>	store into array element as in <code>v1[...] = v2</code>
<code>aload(v1, v2)</code>	load from array element as in <code>v1 = v2[...]</code>
<code>etype(t, et)</code>	array type <code>t</code> has element type <code>et</code>
<code>stype(t1, t2)</code>	type <code>t1</code> is a subtype of type <code>t2</code>
<code>throw(p, v)</code>	program point <code>p</code> throws variable <code>v</code>
<code>catch(t, p, v)</code>	for exceptions at program point <code>p</code> with arg type <code>t</code> , assign arg to <code>v</code>
<code>in(p, m)</code>	program point <code>p</code> is in method <code>m</code>

The analysis defines the following predicates and infers facts of them using the rules in Fig. 2.

- `v_pt(v, c, h, hc)`: variable `v` in context `c` may point to heap abstraction `h` in heap context `hc`
- `f_pt(h1, hc1, f, h2, hc2)`: heap abstraction `h1` in heap context `hc1` may have its field `f` pointing to heap abstraction `h2` in heap context `hc2`
- `r(m, c)`: method `m` is reached in context `c`
- `call(p, c1, m, c2)`: program point `p` in context `c1` calls method `m` in context `c2`
- `assign(v1, c1, v2, c2)`: variable `v1` in context `c1` is assigned the value of `v2` in context `c2`

Each rule gives a direct implication based on the meaning of the predicate. For example, rule (R15) says: if method `m` is reached in context `P(a, b)`, and variable `v` is assigned a new heap abstraction `h` in method `m`, then `v` in context `P(a, b)` may point to `h` in heap context `a`. Note that a context is a pair represented with function symbol `P` since the analysis is 2-call-site sensitive, and that an initial fact `r(main, P(null, null))` can be used to indicate that method `main` is reached in an initial context where the last two calls before calling `main` are `null`.

Following the method above, we first show that the rules are size-bounding for `P` (the only function symbol). (R16), (R17), (R18), (R20), (R21) are trivially size-bounding since they have no function symbols in the conclusion. (R15) is size-bounding since the term of `P` appearing in the conclusion is identical to the one in the hypotheses. (R19) is size-bounding since (i) `P(a, b)` in the conclusion is identical to an occurrence of `P` in the hypotheses, and (ii) for `P(p, a)`, `p` is a program point (i.e., a constant), and therefore cannot have a larger size than `b`, and the size of `P(p, a)` is no more than `P(a, b)`.

Next, we determine which arguments of which predicates are bounded by `count(P)`. These are `v_pt.2` due to (R15) and (R19), `r.2` due to (R19), `call.2` and `call.4` due to (R19), and `assign.2` and `assign.4` due to (R20) and (R21).

Finally, we determine `count(P)`. The only rule whose conclusion contains a term of `P` distinct from the terms in its hypotheses is (R19). The variables of this term of `P` are `p` and `a`. The source of `p` is the third argument of the first hypothesis, therefore its domain size is `#vcall.3`. The source of `a` is the first argument in the terms of `P` in the second and third hypotheses, but this argument of

```

v_pt(v,P(a,b),h,a) ← r(m,P(a,b)), alloc(v,h,m). (R15)
v_pt(v,c,h,hc) ← move(v,v2), v_pt(v2,c,h,hc). (R16)
f_pt(h1,hc1,f,h2,hc2) ← store(v1,f,v2), v_pt(v2,c,h2,hc2), v_pt(v1,c,h1,hc1). (R17)
v_pt(v,c,h,hc) ← load(v,v2,f), v_pt(v2,c,h2,hc2), f_pt(h2,hc2,f,h,hc). (R18)
r(m,P(p,a)), v_pt(t,P(p,a),h,hc), call(p,P(a,b),m2,P(p,a)) ← vcall(v,s,p,m1),
r(m1,P(a,b)), v_pt(v,P(a,b),h,hc), htype(h,ht), lookup(ht,s,m2), this(m2,t). (R19)
assign(v1,c2,v2,c1) ← call(p,c1,m,c2), farg(m,n,v1), aarg(p,n,v2). (R20)
assign(v1,c1,v2,c2) ← call(p,c1,m,c2), aret(p,v1), fret(m,v2). (R21)
v_pt(v,c,h,hc) ← assign(v,c,v2,c2), v_pt(v2,c2,h,hc). (R22)

// decomposed rules:
int23(v1,v2,h1,hc1,h2,hc2) ← v_pt(v2,c,h2,hc2), v_pt(v1,c,h1,hc1). (R17/1)
f_pt(h1,hc1,f,h2,hc2) ← store(v1,f,v2), int23(v1,v2,h1,hc1,h2,hc2). (R17/2)
int24(v,f,c,h2,hc2) ← load(v,v2,f), v_pt(v2,c,h2,hc2). (R18/1)
v_pt(v,c,h,hc) ← int24(v,f,c,h2,hc2), f_pt(h2,hc2,f,h,hc). (R18/2)
int25(s,p,m1,P(a,b),h,hc) ← vcall(v,s,p,m1), v_pt(v,P(a,b),h,hc). (R19/1)
int26(s,p,P(a,b),h,hc) ← int25(s,p,m1,P(a,b),h,hc), r(m1,P(a,b)). (R19/2)
int27(s,p,P(a,b),h,hc,ht,m2) ← int26(s,p,P(a,b),h,hc), lookup(ht,s,m2). (R19/3)
int28(s,p,P(a,b),h,hc,m2) ← int27(s,p,P(a,b),h,hc,ht,m2), htype(h,ht). (R19/4)
r(m,P(p,a)), v_pt(t,P(p,a),h,hc), call(p,P(a,b),m2,P(p,a)) ←
int28(s,p,P(a,b),h,hc,m2), this(m2,t). (R19/5)
int29(c1,m,c2,n,v2) ← call(p,c1,m,c2), aarg(p,n,v2). (R20/1)
assign(v1,c2,v2,c1) ← int29(c1,m,c2,n,v2), farg(m,n,v1). (R20/2)
int30(c1,m,c2,v1) ← call(p,c1,m,c2), aret(p,v1). (R21/1)
assign(v1,c1,v2,c2) ← int30(c1,m,c2,v1), fret(m,v2). (R21/2)

```

Fig. 2. A 2-call-site sensitive analysis with a 1-call-site sensitive heap, and decomposed rules

Table 2. Complexities for a 2-call-site sensitive analysis with a 1-call-site sensitive heap

(R15)	$O(\min(\#r \times \#\text{alloc}.1, 2/3, \#\text{alloc} \times \#r.2/1))$
(R16)	$O(\min(\#\text{move} \times \#v_pt.2, 3, 4/1, \#v_pt \times \#\text{move}.1/2))$
(R17/1)	$O(\#v_pt \times \#v_pt.1, 3, 4/2)$
(R17/2)	$O(\min(\#\text{store} \times \#\text{int23}.3, 4, 5, 6/1, 2, \#\text{int23} \times \#\text{store}.2/1, 3))$
(R18/1)	$O(\min(\#\text{load} \times \#v_pt.2, 3, 4/1, \#v_pt \times \#\text{load}.1, 3/2))$
(R18/2)	$O(\min(\#\text{int24} \times \#\text{f_pt}.4, 5/1, 2, 3, \#\text{f_pt} \times \#\text{int24}.1, 3/2, 4, 5))$
(R19/1)	$O(\min(\#\text{vcall} \times \#v_pt.2, 3, 4/1, \#v_pt \times \#\text{vcall}.2, 3, 4/1))$
(R19/2)	$O(\#\text{int25})$
(R19/3)	$O(\min(\#\text{int26} \times \#\text{lookup}.1, 3/2, \#\text{lookup} \times \#\text{int26}.2, 3, 4, 5/1))$
(R19/4)	$O(\#\text{int27})$
(R19/5)	$O(\min(\#\text{int28} \times \#\text{this}.2/1, \#\text{this} \times \#\text{int28}.1, 2, 3, 4, 5/6))$
(R20/1)	$O(\min(\#\text{call} \times \#\text{aarg}.2, 3/1, \#\text{aarg} \times \#\text{call}.2, 3, 4/1))$
(R20/2)	$O(\min(\#\text{int29} \times \#\text{farg}.3/1, 2, \#\text{farg} \times \#\text{int29}.1, 3, 5/2, 4))$
(R21/1)	$O(\min(\#\text{call} \times \#\text{aret}.2/1, \#\text{aret} \times \#\text{call}.2, 3, 4/1))$
(R21/2)	$O(\min(\#\text{int30} \times \#\text{fret}.2/1, \#\text{fret} \times \#\text{int30}.1, 3, 4/2))$

P is, as just analyzed, only from the third argument of the first hypothesis. Therefore, count(P) is $O((\#\text{vcall}.3)^2)$.

We decompose the rules with our algorithm as shown in Fig. 2, and the calculated complexities are shown in Table 2.

```

must_pt(var,h) ← alloc(_,var,h,_).
must_pt(to,h) ← move(_,to,from), must_pt(from,h).
must_pt(to,h) ← phi(_,to,from1,from2), must_pt(from1,h), must_pt(from2,h).
must_pt(to,h2) ← load(i,to,v,f), must_pt(v,h), f_must_pt(i,h,f,h2).
f_must_pt(i,h,f,h2) ← store(i,v,f,from), must_pt(from,h2), must_pt(v,h).
f_must_pt(j,h,f,h2) ←
  next(_,j), f_must_pt(_,h,f,h2), (forall i: next(i,j) → f_must_pt(i,h,f,h2)),
  not store(j,_,f,_), not vcall(_,_,j,_), not alloc(j,_,h), not alloc(j,_,h2).
f_must_pt(j,h,f,h2) ←
  next(_,j), f_must_pt(_,h,f,h2), (forall i: next(i,j) → f_must_pt(i,h,f,h2)),
  store(j,v,f,_), must_pt(v,h3), h3 != h.

```

Fig. 3. A flow-sensitive must-point-to analysis in Datalog with universal quantification, negation, and inequality

5 Handling rules with universal quantification—applied to flow-sensitive must-point-to analysis

Must-point-to analysis determines the heap abstractions that a pointer variable or expression must refer to, as opposed to may refer to, in all program executions. Flow-sensitive analysis determines analysis results specific to each program point, as opposed to one global result for the program. Therefore flow-sensitive must-point-to analysis can give significantly more certain results that complement flow-insensitive may-point-to alias analysis. This analysis poses two new challenges:

1. The analysis is much more complex, requiring extensions to Datalog with universal quantification and negation.
2. The analysis algorithm is much more sophisticated, requiring new techniques to keep the complexity from increasing.

Specification using Datalog rules with universal quantification and negation. The analysis is specified using seven rules (Smaragdakis and Balatsouras 2015), shown in Fig. 3 (after the changes noted in the third paragraph below). The last two rules are the core of the analysis. The first five rules define `must_pt` and a simple case of `f_must_pt`, where `alloc`, `move`, `load`, and `store` are as in Section 3.2 except with an additional first argument indicating the program point, and `phi` is an instruction for merging the values of two variables. The first five rules are simple Datalog rules with one, two, or three hypotheses; they can be analyzed using the method in Sections 2.1 and 3.1, yielding a time complexity of $O(p \times h^2)$. This section focuses on the two core rules, which are Datalog extended with universal quantification, simple negation, as well as inequality. These two rules are the core of the flow-sensitive analysis because they infer field-must-point-to information for each program point by combining information from all its predecessor points.

The first core rule says that, just after instruction `j`, `h` must point via its field `f` to `h2` if (1) `j` is the next instruction of some instruction, (2) `h` must point via `f` to `h2` just after some instruction, (3) for all instructions `i` just before `j`, `h` must point via `f` to `h2` at `i`, and (4) `j` is not a `store`, `vcall`, or `alloc` instruction that can change the must-point-to information.

The second core rule concludes the same if the same conditions hold except that `j` is a `store` instruction into field `f` of `v`, and `v` must point to a heap abstraction `h3` that is not `h`.

Note that, compared to the original two core rules (Smaragdakis and Balatsouras 2015), we added the first two hypotheses in each rule and moved the conditions about `j` out of the universal

quantification. The two new hypotheses bind the free variables not bound by the universal quantifications, and are necessary for the rules to be correct; without them, the universal quantification returns true when no i satisfies its domain condition $\text{next}(i, j)$, which would lead to f_must_pt to hold for all values of h, f , and $h2$ for any j for which no i satisfies $\text{next}(i, j)$. The conditions about j are moved out because they do not depend on the universally quantified variable i . These also show that the analysis is complex and universal quantification is challenging.

Analysis algorithm for universal quantification and inequality. Despite negation and inequality in the core rules, the set of f_must_pt facts that can be inferred still increases monotonically. Therefore, the set can be computed as a least fixed point as for Datalog. However, if computed straightforwardly, universal quantification adds a linear factor after each f_must_pt fact is added. We show how to compute it, as well as the negation and inequality, incrementally in $O(1)$ time.

Consider the universal quantification, in both core rules:

```
(forall i: next(i, j) → f_must_pt(i, h, f, h2))
```

To compute it efficiently, we maintain the following four auxiliary invariants:

```
prev[j] = {i: (i, j) in next} and
prev_count[j] = #prev[j], for j in next.2

prev_pt[j, h, f, h2] = {i: (i, j) in next, (i, h, f, h2) in f_must_pt} and
prev_pt_count[j, h, f, h2] = #prev_pt[j, h, f, h2], for j in next.2, (h, f, h2) in f_must_pt.2, 3, 4
```

and replace the universal quantification with the following $O(1)$ time test between two aggregate count values:

```
prev_count[j] = prev_pt_count[j, h, f, h2]
```

Variables prev and prev_count for the first two invariants are initialized by iterating over each element (i, j) of input next , adding i to $\text{prev}[j]$ and incrementing $\text{prev_count}[j]$, in a total of $O(\#\text{next})$ time. The next two invariants are maintained incrementally at addition of $(i, h, f, h2)$ to f_must_pt as follows, taking a total of $O(\#\text{next}.2/1 \times \text{f_must_pt})$ time overall all additions:

```
for j in next[i]:
    prev_pt[j, h, f, h2] ∪= {i} // use next to get each next node
    prev_pt_count[j, h, f, h2] += 1 // i is new to h, f, h2 because (i, h, f, h2) is new
    // increment the corresponding count by 1
```

The first core rule now becomes Datalog with simple negations as $O(1)$ time tests, and with the universal quantification as an $O(1)$ time equality test between two counts. Its total time complexity is $O(\#\text{next}.2 \times \#\text{f_must_pt}.2, 3, 4)$.

The second core rule is similar in terms of the universal quantification, but it does not have simple negations but an inequality. We handle the inequality specially, replacing the last two hypotheses on the last line with the following, removing the extra variable $h3$:

```
must_pt[v] - {h} != {} // use must_pt[v] to get heap abstractions that v must point to
```

There is only one value for v in a store instruction, and the above element subtraction and test take $O(1)$ time. So the total time complexity of this rule is again $O(\#\text{next}.2 \times \#\text{f_must_pt}.2, 3, 4)$.

Complexity guarantees. Summing all time complexities together, from initialization, maintaining auxiliary invariants, and using the two resulting rules, yields the total time complexity

$$O(\#\text{next} + \#\text{next}.2/1 \times \#\text{f_must_pt} + \#\text{next}.2 \times \#\text{f_must_pt}.2, 3, 4)$$

$\#\text{next}$ is bounded by $O(p)$, the size of the program. $\text{f_must_pt}.2, 3, 4$ is bounded by the domain sizes of its three arguments $O(h \times f \times h)$. Therefore, the total time complexity is $O(p \times h^2 \times f)$.

6 Additional pointer analyses and summary of complexity analysis results

Besides the 3 pointer analyses discussed, we also studied the remaining 6 analyses in (Smaragdakis and Balatsouras 2015), including 3 in Appendix A; we do not present the rest in detail because they are simpler and do not illustrate additional logic rule features for complexity analysis. There are also other analyses that can be specified using Datalog rules, such as the context-free-language formulation in (Zheng and Rugina 2008). We believe that the reader can follow our method to produce a set of rules and analyze their complexities easily.

In addition to the precise complexities that we calculated, here we also present the worst-case complexities in simpler terms, and provide conditions under which the complexities are linear or quadratic. Table 3 summarizes, for each analysis, the features used, maximum number of hypotheses in the rules, worst-case complexities, and complexities conditioned on constraints on sizes of predicates. We denote an n -call-site sensitive analysis with an m -call-site sensitive heap, as (n, m) -context. The conditions used are as follows, where the conditions on all the EDB predicates are typical for real programs.

$$(C1): O(\#v_pt.2/1)=O(1)$$

$$(C2): O(\#f_pt.3/1,2)=O(1)$$

$$(C3): O(\#lookup.1,3/2)=O(\#this.2/1)=O(\#call.2/1)=O(\#farg.3/1,2)=O(\#fret.2/1)=O(1)$$

$$(C4): O(\#htype.2/1)=O(\#in.2/1)=O(\#t_pt.2/1)=O(\#throw.2/1)=O(1)$$

$$(C5): O(\#must_pt.2/1)=O(1)$$

$$(C6): O(\#f_must_pt.2,3,4)=O(h)$$

(C1) says that each variable may point to a constant number of heap abstractions. (C2) says that each field of each variable may point to a constant number of heap abstractions. The rest are similar.

Complexities of pointer analysis with constraints on the sizes of program parameters have been studied. One can obtain the complexities achieved in such studies using our method, and substituting the relevant complexity parameters in our analysis with the constraints. In (Sridharan and Fink 2009), the authors present an algorithm for Andersen’s analysis, which runs in $O((v+h)^2)$ time for k -sparse programs, where v is the number of variables and h is the number of heap abstractions. The definition of k -sparse programs has two constraints: For our rules in Figure 1, the first constraint implies $O(\#store.2,3/1) = O(\#load.1,3/2) = O(1)$, and the second constraint implies that $\#move+\#int1+\#int2 \leq O(v+h)$. Substituting these constraints in our complexity analysis in Table 1, we obtain $O((v+h) \times (\#v_pt.2/1 + \#f_pt.3/1,2))$, which is in the worst case $O((v+h) \times h)$. Thus, we obtain a better and more precise complexity than (Sridharan and Fink 2009).

For Andersen OO analysis, interprocedural analysis, and $(0,1)$ -context-sensitive analysis, the worst-case complexities when parametrized by only program size, n , are known—an upper bound of $O(n^3)$ for the first two (Andersen 1994) and $O(n^5)$ for the third (Wilson and Lam 1995). However, to our knowledge, we present more precise complexities for these analyses for the first time, and for the other analyses, we present complexities for the first time. Our complexity results are improvements since they are tighter than known worst-case complexities, and when our fine-grained analyses are used, the running time of the analyses can be better understood.

Table 3. Summary of complexities. Recall p , h , and f are number of program points, heap abstractions, and fields, respectively.

Analysis	Features	#Hypo.	Worst-case	Conditional
Andersen OO (Sec. 3)	Pure Datalog	3	$O(p \times h^2)$	$O(p \times h)$ if (C1), $O(p)$ if (C1) and (C2)
Interprocedural (Sec. A.1) ²	Pure Datalog	6	$O(p^2 \times h)$	$O(p)$ if (C1), (C2) and (C3)
Arrays (Sec. A.2) ²	Pure Datalog	7	$O(p^2 \times h)$	$O(p)$ if (C1), (C2) and (C3)
Exceptions (Sec. A.3) ²	Negation	5	$O(p^2 \times h)$	$O(p)$ if (C1) and (C3)
Reflection	Pure Datalog	9	$O(p^3 \times h)$	$O(p \times h)$ obtainable ³
(0,1)-context	Function symbols	6	$O(p^3 \times h^2)$	–
(1,1)-context	Function symbols	6	$O(p^4 \times h^2)$	–
(2,1)-context (Sec. 4)	Function symbols	6	$O(p^5 \times h^2)$	–
Flow must (Sec. 5)	Univ. quant., negation, inequality	8	$O(p \times h^2 \times f)$	$O(p \times h)$ if (C5) and (C6)

7 Related work and conclusion

We discuss related work on Datalog evaluation, applications of rules to pointer analysis, precise complexities for pointer analyses, and directions for future work.

Evaluation of Datalog has been studied for a long time (Ceri et al. 1990). Optimal bottom-up evaluation of Datalog rules with complexity guarantees is first given in (Liu and Stoller 2009), but no algorithm is given for decomposing rules except for trying all decompositions. We build on this method for evaluating Datalog rules and calculating complexities, but extend it to handle rules with many hypotheses and other Datalog extensions. Our new algorithm and method are able to obtain new or more precise complexities compared with the best previous complexities, as discussed in Section 6.

Formulation of various static program analyses as rules has been studied. In particular, Andersen’s pointer analysis (Andersen 1994) was formulated as deductive rules in (Heintze and Tardieu 2001), and given as logic rules in (Saha and Ramakrishnan 2005). Andersen’s analysis with many flavors was given in a recent survey (Smaragdakis and Balatsouras 2015), on which we base our study. The fact that the time complexity of Andersen’s analysis is worst-case cubic has been known since the original introduction. (Sridharan and Fink 2009) notes that typical behavior is different than worst case and proves that under certain conditions the analysis is quadratic. We give the precise time complexities for Andersen’s analysis for an object-oriented language, and show also precise conditions, additional to existing literature, under which the complexities are linear or quadratic directly as results of our complexity analyses. We also obtain precise complexities for pointer analyses for additional language features, and provide methods for handling extensions to Datalog such as function symbols when such extensions are necessary to implement the analyses.

Must-point-to analyses are more complex, and rules modeling the analyses involve universal quantification and inequality beyond pure Datalog. (Hind et al. 1999) gives an $O(n^5)$ algorithm for a must-alias analysis (closely related to must-point-to analyses); using methods that are also

² In Appendix A.

³ The conditions to obtain this complexity involve EDB predicates not discussed in this paper.

employed in optimal bottom-up evaluation of Datalog, (Goyal 2005) improves this complexity to $O(n^3)$. In this paper, we show how to handle must-point-to analysis expressed using rules extended with universal quantification and a special inequality, and provide precise complexity analyses for our efficient implementation. Transforming quantifications into aggregate queries such as counts has been used in other applications, e.g., distributed algorithms (Liu et al. 2012), but how to handle inequality in general is a subject for future study.

Future directions include analyzing and optimizing the space complexity of pointer analyses, especially to remove unnecessary intermediate predicates introduced for rules with many hypotheses, and optimization of demand-driven pointer analysis via queries, e.g., by using the methods of (Tekle and Liu 2010) and (Tekle and Liu 2011).

References

- Lars Ole Andersen. 1994. *Program Analysis and Specialization for the C Programming Language*. Ph.D. Dissertation. DIKU, University of Copenhagen, Copenhagen, Denmark.
- Dzintars Avots, Michael Dalton, V. Benjamin Livshits, and Monica S. Lam. 2005. Improving Software Security with a C Pointer Analysis. In *Proc. of the 27th Intl. Conf. on Software Engineering*. 332–341.
- Stefano Ceri, Georg Gottlob, and Letizia Tanca. 1990. *Logic Programming and Databases*. Springer.
- Amer Diwan, Kathryn S. McKinley, and J. Eliot B. Moss. 1998. Type-Based Alias Analysis. In *Proc. of the ACM SIGPLAN Conf. on Programming Language Design and Implementation*. 106–117.
- Rakesh Ghiya and Laurie J. Hendren. 1998. Putting Pointer Analysis to Work. In *Proc. of the 25th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*. 121–133.
- Rakesh Ghiya, Daniel M. Lavery, and David C. Sehr. 2001. On the Importance of Points-to Analysis and Other Memory Disambiguation Methods for C Programs. In *Proc. of the 2001 ACM SIGPLAN Conf. on Programming Language Design and Implementation*. 47–58.
- Michael Gorbovitski, Yanhong A. Liu, Scott D. Stoller, Tom Rothamel, and Tuncay Tekle. 2010. Alias Analysis for Optimization of Dynamic Languages. In *Proceedings of the 6th Symposium on Dynamic Languages*. ACM Press, 27–42.
- Deepak Goyal. 2005. Transformational Derivation of an Improved Alias Analysis Algorithm. *Higher-Order and Symbolic Computation* 18, 1-2 (2005), 15–49.
- Ben Hardekopf and Calvin Lin. 2007. The Ant and the Grasshopper: Fast and Accurate Pointer Analysis for Millions of Lines of Code. In *Proc. of the ACM SIGPLAN 2007 Conf. on Programming Language Design and Implementation*. 290–299.
- Nevin Heintze and Olivier Tardieu. 2001. Ultra-fast Aliasing Analysis using CLA: A Million Lines of C Code in a Second. In *Proc. of the 2001 ACM SIGPLAN Conf. on Programming Language Design and Implementation*. 254–263.
- Michael Hind. 2001. Pointer Analysis: Haven't We Solved This Problem Yet?. In *Proc. of the 2001 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis For Software Tools and Engineering*. 54–61.
- Michael Hind, Michael G. Burke, Paul R. Carini, and Jong-Deok Choi. 1999. Interprocedural Pointer Alias Analysis. *ACM Trans. Program. Lang. Syst.* 21, 4 (1999), 848–894.
- Michael Hind and Anthony Pioli. 2000. Which Pointer Analysis Should I Use?. In *Proc. of the 2000 ACM SIGSOFT Intl. Symp. on Software Testing and Analysis*. 113–123.
- William Landi. 1992. Undecidability of Static Analysis. *ACM Letters on Programming Languages and Systems* 1, 4 (1992), 323–337.
- Yanhong A. Liu and Scott D. Stoller. 2009. From Datalog Rules to Efficient Programs with Time and Space Guarantees. *ACM Trans. Program. Lang. Syst.* 31, 6, Article 21 (Aug. 2009), 38 pages.
- Yanhong A. Liu, Scott D. Stoller, Bo Lin, and Michael Gorbovitski. 2012. From Clarity to Efficiency for Distributed Algorithms. In *Proc. of the 27th ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages and Applications*. 395–410.

- Ana Milanova, Atanas Rountev, and Barbara G. Ryder. 2005. Parameterized Object Sensitivity for Points-to Analysis for Java. *ACM Trans. Softw. Eng. Methodol.* 14, 1 (2005), 1–41.
- David J. Pearce, Paul H. J. Kelly, and Chris Hankin. 2007. Efficient Field-Sensitive Pointer Analysis of C. *ACM Trans. Program. Lang. Syst.* 30, 1, Article 4 (2007).
- G. Ramalingam. 1994. The Undecidability of Aliasing. *ACM Trans. Program. Lang. Syst.* 16, 5 (1994), 1467–1471.
- Shmuel Sagiv, Thomas W. Reps, and Reinhard Wilhelm. 1998. Solving Shape-Analysis Problems in Languages with Destructive Updating. *ACM Trans. Program. Lang. Syst.* 20, 1 (1998), 1–50.
- Diptikalyan Saha and C. R. Ramakrishnan. 2005. Incremental and Demand-Driven Points-to Analysis Using Logic Programming. In *Proc. of the 7th Intl. ACM SIGPLAN Conf. on Principles and Practice of Declarative Programming*. 117–128.
- Danny De Schreye and Stefaan Decorte. 1994. Termination of Logic Programs: The Never-Ending Story. *The Journal of Logic Programming* 19/20 (1994), 199–260.
- Patricia G. Selinger, Morton M. Astrahan, Donald D. Chamberlin, Raymond A. Lorie, and Thomas G. Price. 1979. Access Path Selection in a Relational Database Management System. In *Proc. of the 1979 ACM SIGMOD Intl. Conf. on Management of Data*. 23–34.
- Marc Shapiro and Susan Horwitz. 1997. The Effects of the Precision of Pointer Analysis. In *Proc. of the 4th Intl. Symp. on Static Analysis*. 16–34.
- Micha Sharir and Amir Pnueli. 1981. Two Approaches to Interprocedural Data Flow Analysis. In *Program Flow Analysis: Theory and Applications*. Prentice-Hall, Chapter 7, 189–233.
- Olin Grigsby Shivers. 1991. *Control-flow Analysis of Higher-order Languages of Taming Lambda*. Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA, USA.
- Yannis Smaragdakis and George Balatsouras. 2015. Pointer Analysis. *Foundations and Trends in Programming Languages* 2, 1 (2015), 1–69.
- Yannis Smaragdakis, Martin Bravenboer, and Ondrej Lhoták. 2011. Pick Your Contexts Well: Understanding Object-Sensitivity. In *Proc. of the 38th Symp. on Principles of Programming Languages*. 17–30.
- Manu Sridharan and Stephen J. Fink. 2009. The Complexity of Andersen’s Analysis in Practice. In *Proc. of the 16th Intl. Symp. on Static Analysis*. 205–221.
- Bjarne Steensgaard. 1996. Points-to Analysis in Almost Linear Time. In *Conf. Record of the 23rd ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*. 32–41.
- Michael Steinbrunn, Guido Moerkotte, and Alfons Kemper. 1997. Heuristic and Randomized Optimization for the Join Ordering Problem. *The VLDB Journal* 6, 3 (1997), 191–208.
- K. Tuncay Tekle and Yanhong A. Liu. 2010. Precise Complexity Analysis for Efficient Datalog Queries. In *Proc. of the 12th Intl. ACM SIGPLAN Conf. on Principles and Practice of Declarative Programming*. 35–44.
- K. Tuncay Tekle and Yanhong A. Liu. 2011. More Efficient Datalog Queries: Subsumptive Tabling Beats Magic Sets. In *Proc. of the 2011 ACM SIGMOD Intl. Conf. on Management of Data*. 661–672.
- Robert P. Wilson and Monica S. Lam. 1995. Efficient Context-Sensitive Pointer Analysis for C Programs. In *Proc. of the ACM SIGPLAN’95 Conf. on Programming Language Design and Implementation*. 1–12.
- Xin Zheng and Radu Rugina. 2008. Demand-Driven Alias Analysis for C. In *Proc. of the 35th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*. 197–208.

$v_pt(v,h) \leftarrow \text{alloc}(v,h,m), r(m).$	(R1*)
$r(m2), v_pt(t,h), \text{call}(p,m2) \leftarrow \text{vcall}(v,s,p,m1), r(m1), v_pt(v,h),$	
$\text{hType}(h,ht), \text{lookup}(ht,s,m2), \text{this}(m2,t).$	(R5)
$\text{assign}(v1,v2) \leftarrow \text{call}(p,m), \text{farg}(m,n,v1), \text{aarg}(p,n,v2).$	(R6)
$\text{assign}(v1,v2) \leftarrow \text{call}(p,m), \text{aret}(p,v1), \text{fret}(m,v2).$	(R7)
$v_pt(v1,h) \leftarrow \text{assign}(v1,v2), v_pt(v2,h).$	(R8)
$\text{int3}(v,s,p) \leftarrow \text{vcall}(v,s,p,m1), r(m1).$	(R5/1)
$\text{int4}(s,p,h) \leftarrow \text{int3}(v,s,p), v_pt(v,h).$	(R5/2)
$\text{int5}(p,h,ht,m2) \leftarrow \text{int4}(s,p,h), \text{lookup}(ht,s,m2).$	(R5/3)
$\text{int6}(p,h,m2) \leftarrow \text{int5}(p,h,ht,m2), \text{hType}(h,ht).$	(R5/4)
$r(m2), v_pt(t,h), \text{call}(p,m2) \leftarrow \text{int6}(p,h,m2), \text{this}(m2,t).$	(R5/5)
$\text{int7}(m,n,v2) \leftarrow \text{call}(p,m), \text{aarg}(p,n,v2).$	(R6/1)
$\text{assign}(v1,v2) \leftarrow \text{int7}(m,n,v2), \text{farg}(m,n,v1).$	(R6/2)
$\text{int8}(m,v1) \leftarrow \text{call}(p,m), \text{aret}(p,v1).$	(R7/1)
$\text{assign}(v1,v2) \leftarrow \text{int8}(m,v1), \text{fret}(m,v2).$	(R7/2)

Fig. A 1. Additional and redefined rules for an interprocedural analysis, and decomposed rules

Appendix A Pointer analyses for advanced features and their complexities

Our method for handling many hypotheses applies also to pointer analyses for other advanced language features. We present different analyses (Smaragdakis and Balatsouras 2015) built on the base analysis in Fig. 1. For each analysis, we apply our algorithm for rule decomposition as described before, including for rules with simple negation, and show the complexity calculation.

A more precise pointer analysis can be obtained by making the analysis interprocedural. The rules for an interprocedural version of Andersen’s analysis are more complex, and contain up to 6 hypotheses. We present the rule decompositions and calculated complexities for this analysis in Section A.1.

Another feature to add is to analyze arrays. The analysis includes using a rule with 7 hypotheses. We present the rule decompositions and calculated complexities for analyzing arrays in Section A.2.

We also show an analysis of exceptions. This analysis uses a simple form of negation where only some EDB hypotheses are negated. Such negation is handled during evaluation by simply testing the truth value for the negated hypotheses. We describe the handling of negation during evaluation and rule decomposition, and then present the rule decompositions and calculated complexities for analyzing exceptions in Section A.3.

A.1 An interprocedural may-point-to analysis

In addition to the IDB predicates for the base analysis presented in Section 3.2, the following IDB predicates are used for interprocedural analysis.

- $r(m)$: Method m is reached.
- $\text{call}(p,m)$: Statement at program point p calls method m .
- $\text{assign}(v1,v2)$: $v1$ is assigned $v2$ either as part of argument or return passing.

These IDB predicates are defined with the rules in Fig. A 1. In addition, (R1) is redefined with (R1*) to be more precise, and more precise definitions for v_pt are added. Since (R5), (R6) and (R7) have more than two hypotheses, we decompose them using our algorithm resulting in the rules also shown in Fig. A 1, and we calculate the complexities are shown in Table A 1.

Table A 1. Complexities for an interprocedural point-to analysis

Rule	Complexity
(R1*)	$O(\#\text{alloc})$
(R5/1)	$O(\#\text{vcall})$
(R5/2)	$O(\min(\#\text{int3} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{int3}.2,3/1))$
(R5/3)	$O(\min(\#\text{int4} \times \#\text{lookup}.1,3/2, \#\text{lookup} \times \#\text{int4}.2,3/1))$
(R5/4)	$O(\#\text{int5})$
(R5/5)	$O(\min(\#\text{int6} \times \#\text{this}.2/1, \#\text{this} \times \#\text{int6}.1,2/3))$
(R6/1)	$O(\min(\#\text{call} \times \#\text{aarg}.2,3/1, \#\text{aarg} \times \#\text{call}.2/1))$
(R6/2)	$O(\min(\#\text{int7} \times \#\text{farg}.3/1,2, \#\text{farg} \times \#\text{int7}.3/1,2))$
(R7/1)	$O(\min(\#\text{call} \times \#\text{aret}.2/1, \#\text{aret} \times \#\text{call}.2/1))$
(R7/2)	$O(\min(\#\text{int8} \times \#\text{fret}.2/1, \#\text{fret} \times \#\text{int8}.2/1))$
(R8)	$O(\min(\#\text{assign} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{assign}.1/2))$

$\text{a_pt}(h1,h2) \leftarrow \text{astore}(v1,v2), \text{v_pt}(v1,h1), \text{v_pt}(v2,h2), \text{htype}(h1,ht1),$ $\text{htype}(h2,ht2), \text{etype}(ht2,et), \text{stype}(ht1,et).$	(R9)
$\text{v_pt}(v,h) \leftarrow \text{aload}(v,v2), \text{v_pt}(v2,h2), \text{a_pt}(h2,h).$	(R10)
$\text{int9}(v2,h1) \leftarrow \text{astore}(v1,v2), \text{v_pt}(v1,h1).$	(R9/1)
$\text{int10}(h1,h2) \leftarrow \text{int9}(v2,h1), \text{v_pt}(v2,h2).$	(R9/2)
$\text{int11}(ht1,ht2) \leftarrow \text{etype}(ht2,et), \text{stype}(ht1,et).$	(R9/3)
$\text{int12}(h1,ht2) \leftarrow \text{htype}(h1,ht1), \text{int11}(ht1,ht2).$	(R9/4)
$\text{int13}(h1,h2) \leftarrow \text{int12}(h1,ht2), \text{htype}(h2,ht2).$	(R9/5)
$\text{a_pt}(h1,h2) \leftarrow \text{int10}(h1,h2), \text{int13}(h1,h2).$	(R9/6)
$\text{int14}(v,h2) \leftarrow \text{aload}(v,v2), \text{v_pt}(v2,h2).$	(R10/1)
$\text{v_pt}(v,h) \leftarrow \text{int14}(v,h2), \text{a_pt}(h2,h).$	(R10/2)

Fig. A 2. Additional rules for analyzing arrays, and decomposed rules

A.2 May-point-to analysis for arrays

Analysis of arrays is done in an *array-insensitive* manner, so that the analysis is performed on arrays but different array indices are not distinguished. A new IDB predicate $\text{a_pt}(a,h)$ is defined, indicating that elements of array a may point to heap abstraction h , and a new rule is added for v_pt as shown in Fig. A 2. In the figure, we show the decomposition of these rules using our algorithm. We calculate the complexities from these rules as shown in Table A 2.

Table A 2. Complexities incurred by analyzing arrays

Rule	Complexity
(R9/1)	$O(\min(\#\text{astore} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{astore}.2/1))$
(R9/2)	$O(\min(\#\text{int9} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{int9}.2/1))$
(R9/3)	$O(\min(\#\text{etype} \times \#\text{stype}.1/2, \#\text{stype} \times \#\text{etype}.1/2))$
(R9/4)	$O(\min(\#\text{htype} \times \#\text{int11}.2/1, \#\text{int11} \times \#\text{htype}.1/2))$
(R9/5)	$O(\min(\#\text{int12} \times \#\text{htype}.1/2, \#\text{htype} \times \#\text{int12}.1/2))$
(R9/6)	$O(\#\text{int10})$
(R10/1)	$O(\min(\#\text{aload} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{aload}.1/2))$
(R10/2)	$O(\min(\#\text{int14} \times \#\text{a_pt}.2/1, \#\text{a_pt} \times \#\text{int14}.1/2))$

$t_pt(m,h) \leftarrow in(p,m), throw(p,v), v_pt(v,h), htype(h,ht), not\ catch(ht,p,_)$.	(R11)
$t_pt(m,h) \leftarrow in(p,m), call(p,m2), t_pt(m2,h), htype(h,ht), not\ catch(ht,p,_)$.	(R12)
$v_pt(v2,h) \leftarrow throw(p,v), v_pt(v,h), htype(h,ht), catch(ht,p,v2)$.	(R13)
$v_pt(v,h) \leftarrow call(p,m), t_pt(m,h), htype(h,ht), catch(ht,p,v)$.	(R14)
$int15(p,h) \leftarrow throw(p,v), v_pt(v,h)$.	(R11/1)
$int16(p,h) \leftarrow int15(p,h), htype(h,ht), not\ catch(ht,p,_)$.	(R11/2)
$t_pt(m,h) \leftarrow in(p,m), int16(p,h)$.	(R11/3)
$int17(p,h) \leftarrow call(p,m2), t_pt(m2,h)$.	(R12/1)
$int18(p,h) \leftarrow int17(p,h), htype(h,ht), not\ catch(ht,p,_)$.	(R12/2)
$t_pt(m,h) \leftarrow in(p,m), int18(p,h)$.	(R12/3)
$int19(v,ht,v2) \leftarrow throw(p,v), catch(ht,p,v2)$.	(R13/1)
$int20(h,ht,v2) \leftarrow v_pt(v,h), int19(v,ht,v2)$.	(R13/2)
$v_pt(v2,h) \leftarrow int20(h,ht,v2), htype(h,ht)$.	(R13/3)
$int21(m,ht,v) \leftarrow call(p,m), catch(ht,p,v)$.	(R14/1)
$int22(h,ht,v) \leftarrow int21(m,ht,v), t_pt(m,h)$.	(R14/2)
$v_pt(v,h) \leftarrow int22(h,ht,v), htype(h,ht)$.	(R14/3)

Fig. A 3. Additional rules for analyzing exceptions, and decomposed rules

A.3 Handling negation—applied to for may-point-to analysis for exceptions

A hypothesis can be *negated*, indicated by a preceding `not`. Given a rule with negated hypotheses, and values for the variables, the rule infers the conclusion as a fact if all non-negated hypotheses are true and all negated hypotheses are false for the given values of variables. We say that rules have *simple negation* if the only negated hypotheses in the rules are EDB hypotheses; this negation is simple because it can be performed as a check after all non-negated hypotheses are processed. Therefore, a rule only needs decomposition if there are more than two non-negated hypotheses.

In the case when a negated hypothesis appears in a rule that needs to be decomposed, we apply our algorithm as described before without considering the negated hypotheses, and if at any point all the essential variables of a negated hypotheses appear in an intermediate rule, then we add the negated hypotheses to the rule. A variable of a hypothesis is *essential* in a rule if it appears in any other hypothesis or conclusion; we name non-essential variables with an underscore (`_`) in the rules. After this decomposition, during bottom-up evaluation of a rule with negated hypotheses, for each value combination of variables of non-negated hypotheses that make the hypotheses true, we check if the negated hypotheses hold for the given values of the essential variables and any value of the non-essential variables. If none of the negated hypotheses hold, then we add the conclusion as a fact. This does not change the complexity of the rule since such checks can be performed in constant time.

Exception analysis. Rules are added to allow analysis of the exception flow of a program as shown in Fig. A 3. The rules define a new IDB predicate $t_pt(m,h)$ capturing what heap abstractions h a method m can throw to its catchers for an exception. Rules are also added v_pt for analyzing the exceptions. The rules contain simple negation since `catch` is an EDB predicate. We show the decomposition of the rules as well in Fig. A 3, and the complexities calculated with our method in Table A 3.

Table A 3. Complexities incurred by analyzing exceptions

Rule	Complexity
(R11/1)	$O(\min(\#\text{throw} \times \#\text{v_pt}.2/1, \#\text{v_pt} \times \#\text{throw}.1/2))$
(R11/2)	$O(\min(\#\text{int15} \times \#\text{htype}.2/1, \#\text{htype} \times \#\text{int15}.1/2))$
(R11/3)	$O(\min(\#\text{in} \times \#\text{int16}.2/1, \#\text{int16} \times \#\text{in}.2/1))$
(R12/1)	$O(\min(\#\text{call} \times \#\text{t_pt}.2/1, \#\text{t_pt} \times \#\text{call}.2/1))$
(R12/2)	$O(\min(\#\text{int17} \times \#\text{htype}.2/1, \#\text{htype} \times \#\text{int17}.1/2))$
(R12/3)	$O(\min(\#\text{in} \times \#\text{int18}.2/1, \#\text{int18} \times \#\text{in}.2/1))$
(R13/1)	$O(\min(\#\text{throw} \times \#\text{catch}.1,3/2, \#\text{catch} \times \#\text{throw}.2/1))$
(R13/2)	$O(\min(\#\text{v_pt} \times \#\text{int19}.2,3/1, \#\text{int19} \times \#\text{v_pt}.2/1))$
(R13/3)	$O(\#\text{int20})$
(R14/1)	$O(\min(\#\text{call} \times \#\text{catch}.1,3/2, \#\text{catch} \times \#\text{call}.2/1))$
(R14/2)	$O(\min(\#\text{int21} \times \#\text{t_pt}.2/1, \#\text{t_pt} \times \#\text{int21}.2,3/1))$
(R14/3)	$O(\#\text{int22})$