

1 Functional Programming with Datalog

2 André Pacak

3 JGU Mainz, Germany

4 Sebastian Erdweg

5 JGU Mainz, Germany

6 Abstract

7 Datalog is a carefully restricted logic programming language. What makes Datalog attractive is its
8 declarative fixpoint semantics: Datalog queries consist of simple Horn clauses, yet Datalog solvers
9 efficiently compute all derivable tuples even for recursive queries. However, as we argue in this
10 paper, Datalog is ill-suited as a programming language and Datalog programs are hard to write
11 and maintain. We propose a “new” frontend for Datalog: functional programming with sets called
12 *functional IncA*. While programmers write recursive functions over algebraic data types and sets,
13 we transparently translate all code to Datalog relations. However, we retain Datalog’s strengths:
14 Functions that generate sets can encode arbitrary relations and mutually recursive functions have
15 fixpoint semantics. We also ensure that the generated Datalog program terminates whenever the
16 original functional program terminates, so that we can apply off-the-shelf bottom-up Datalog
17 solvers. We demonstrate the versatility and ease of use of functional IncA by implementing a type
18 checker, a program transformation, an interpreter of the untyped lambda calculus, two data-flow
19 analyses, and clone detection of Java bytecode.

20 **2012 ACM Subject Classification** Software and its engineering → Software notations and tools

21 **Keywords and phrases** Datalog, functional programming, demand transformation

22 **Digital Object Identifier** 10.4230/LIPIcs...

23 1 Introduction

24 Datalog is a carefully restricted logic programming language that has seen a surge in popularity
25 in recent years. Originally, Datalog was conceived as a database query language that operates
26 on finite sets only [15], so that all queries are guaranteed to terminate. Nowadays, Datalog is
27 being used in a wide array of applications [12], from program analysis [10, 13, 23] to network
28 monitoring [1] and distributed computing [2, 3]. What makes Datalog so popular is that (i)
29 there are highly efficient and scalable implementations available and (ii) Datalog programs
30 are considered declarative. We argue that the latter is partly a misconception: Datalog’s
31 semantics is declarative, but Datalog’s frontend is not.

32 Datalog is often primed as being declarative. This can be surprising given that a Datalog
33 program consists of simple Horn clauses ($a_0 :- a_1, \dots, a_n$), where a_0 holds if a_1 through a_n
34 hold. In Datalog, a_0 is called the *head* of the rule and a_1, \dots, a_n form the *body* of the rule.
35 Both head and body consist of *atoms* a , which are of the form $R(t_1, \dots, t_n)$ for some relation
36 R and terms t . A Datalog solver computes the least fixpoint of the Horn clauses such that
37 the relations R contain all derivable ground tuples, called *facts* in Datalog. In the initial
38 fixpoint iteration, the semantics collects all rule heads a_0 that have no precondition. In
39 subsequent fixpoint iterations, the semantics collects all facts that can be derived by applying
40 rules to previously derived facts. When terms range over finite sets, this fixpoint iteration
41 terminates in finitely many steps. We concur that Datalog has a declarative semantics,
42 because programmers do not need to think about *how* the derivable facts are computed.

43 The problem of Datalog is its frontend: It is ill-suited as a programming language and
44 not declarative. Consider we want to construct control-flow graphs as a basis for program
45 analysis. Figure 1 shows a functional program and a Datalog program that construct the
46 control-flow graphs for the While language. The functional program uses pattern matching



© André Pacak and Sebastian Erdweg;

licensed under Creative Commons License CC-BY 4.0

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

```

// Functional programming
def flow(stm: Stm): Set[(Stm, Stm)] = stm match {
  case Assign(x, a) => {}
  case Sequence(s1, s2) => flow(s1) ++ flow(s2) ++ {(l1, init(s2)) | l1 in final(s1)}
  case If(c, s1, s2) => c match {
    case True() => flow(s1) ++ {(stm, init(s1))}
    case False() => flow(s2) ++ {(stm, init(s2))}
    case _ => flow(s1) ++ flow(s2) ++ {(stm, init(s1)), (stm, init(s2))} }
  case While(c, s) => flow(s) ++ {(stm, init(s))} ++ {(l,stm) | l in final(s) }

// Datalog
flow(Stm, From, To) :- sequence(Stm, Stm1, _), flow(Stm1, From, To).
flow(Stm, From, To) :- sequence(Stm, _, Stm2), flow(Stm2, From, To).
flow(Stm, From, To) :- sequence(Stm, Stm1, Stm2), final(Stm1, From), init(Stm2, To).
flow(Stm, From, To) :- if(Stm, C, Stm1, _), true(C), flow(Stm1, From, To).
flow(Stm, From, To) :- if(Stm, C, Stm1, _), true(C), init(Stm1, To).
flow(Stm, From, To) :- if(Stm, C, _, Stm2, _), false(C), flow(Stm2, From, To).
flow(Stm, From, To) :- if(Stm, C, _, Stm2, _), false(C), init(Stm2, To).
flow(Stm, From, To) :- if(Stm, C, Stm1, _), not true(C), not false(C), flow(Stm1,From,To).
flow(Stm, From, To) :- if(Stm, C, Stm1, _), not true(C), not false(C), init(Stm1,To).
flow(Stm, From, To) :- if(Stm, C,_,Stm2,_) , not true(C), not false(C), flow(Stm2,From,To).
flow(Stm, From, To) :- if(Stm, C,_,Stm2,_) , not true(C), not false(C), init(Stm2,To).
flow(Stm, From, To) :- while(Stm, _, Stm1), flow(Stm1, From, To).
flow(Stm, From, To) :- while(Stm, _, Stm1), init(Stm1, To).
flow(Stm, From, To) :- while(Stm, _, Stm1), final(Stm1, From).

```

■ **Figure 1** Constructing control-flow graphs using functional programming and Datalog.

47 and set-comprehensions to compute sets of edges similar to [16], whereas the Datalog program
 48 provides rules to constrain the logic variables `From` and `To`. The most prominent problem with
 49 Datalog in this example is the lack of structured programming and the duplication of atoms,
 50 especially for *if*-statements: we must query relation `if` 8 times, relation `true` 6 times, and
 51 relation `false` 6 times. Such Datalog code is hard to write and maintain. Another problem
 52 with programming Datalog is that rules must be *range-restricted*: Each variable in the head
 53 of a rule must be bound in the body of the rule. This restriction ensures relations can be
 54 computed using Datalog’s least fixpoint semantics. For example, the increment relation
 55 `inc(X, Y) :- Y=X+1` would be correctly rejected by Datalog solvers such as Soufflé [20], because
 56 `x` is not bound in the rule’s body. Datalog programmers need to work around this restriction.

57 So why would programmers want to use Datalog anyways instead of functional program-
 58 ming? Because of Datalog’s declarative fixpoint semantics, which makes it easy to process
 59 cyclic data structures such as our control-flow graphs from above. For example, we can
 60 compute the transitive control-flow reachability with two simple rules:

```

61 flowTrans(Prog, From, To) :- flow(Prog, From, To).
62 flowTrans(Prog, From, To) :- flow(Prog, From, Inter), flowTrans(Prog, Inter, To).
63

```

65 What is remarkable is that we do not have to implement a termination condition and or
 66 detect when the relations are stable; Datalog takes care of that. This is why Datalog is a
 67 popular implementation language for data-flow analyses that propagate information along
 68 the control-flow graph until reaching a fixpoint [10], despite the shortcomings of its frontend.

69 In this paper, we design a functional programming language with fixpoint semantics
 70 and propose it as a “new” Datalog frontend: *functional InCA*. In particular, we show how
 71 functional programs with first-order functions and recursive algebraic data types can be
 72 faithfully translated to negation-free Datalog. A key idea of our approach is to systematically
 73 track the *demand* on functions: Which inputs must a function be run on to obtain the

```

def entry_var(stm: Stm, prog: Stm, x: String): Val =
  fold(BotVal(), joinVal, {exit_var(pred, prog, x) | (pred,stm) in flow(prog)})
def exit_var(stm: Stm, prog: Stm, x: String): Val = stm match {
  case Assign(y, exp) =>
    if (x == y) aeval(exp, stm, prog)
    else entry_var(stm, prog, x)
  case ... }
def aeval(exp: Exp, node: Stm, prog: Stm): Val = exp match {
  case Var(x) => entry_var(node, prog, x)
  case ... }

```

■ **Figure 2** A data-flow analysis using functional programming with fixpoint semantics.

74 computation’s final result. Since terminating functional programs only consider finitely many
 75 inputs, we can track these inputs in Datalog relations. For programs without algebraic data
 76 types, we can adopt a standard demand transformation [25]. However, for algebraic data
 77 types we need to carefully instrument the demand transformation to encode constructors and
 78 selectors through finite relations. Our translation preserves the semantics of the functional
 79 program and, in particular, the resulting Datalog program terminates whenever the functional
 80 program does. The translation targets Datalog with base types and operations on such types
 81 such as integers, float, strings, booleans as well as algebraic data types. The most common
 82 Datalog dialects such as Soufflé, IncA, Flix and Formulog all support these types. Whenever
 83 we reference Datalog we mean Datalog with the extensions listed above.

84 Functional IncA replaces Datalog’s logic programming frontend, but we retain Datalog’s
 85 key advantages: relations with fixpoint semantics. Specifically, we extend functional IncA
 86 with set types and set operations (comprehensions, union, and folds) such that programmers
 87 can describe and aggregate over relations. For example, Figure 2 shows a data-flow analysis
 88 implemented in functional IncA. The analysis queries the control-flow graph `flow` and propa-
 89 gates information about the value of variables (abstracted as intervals) along the control-flow
 90 graph. Note that `entry_var`, `exit_var`, and `aeval` are mutually recursive and that there is no
 91 termination condition (the program diverges under standard functional semantics). Despite
 92 using functional programming as a frontend, all code compiles to Datalog rules, which is a key
 93 advantage of our approach for two reasons. First, programmers can rely on the declarative
 94 Datalog semantics to find the least fixpoint. Second, we can use any existing Datalog solver
 95 to run the program, whereas prior Datalog dialects usually require a custom Datalog solver.

96 We have implemented functional IncA as part of the incremental Datalog framework
 97 IncA [23]. We translate functional IncA into a Datalog IR and provide two backends:
 98 one targeting IncA directly, the other targeting Soufflé [20]. While the IncA backend
 99 provides incremental re-evaluation after input changes, the Soufflé backend provides better
 100 non-incremental performance. The choice of the backend is transparent for the user of
 101 the frontend, except that Soufflé does not support user-defined aggregations. We have
 102 implemented three case studies using functional IncA. First, we implemented a type checker
 103 for the simply-typed lambda calculus, a type-erasure transformation for the same, and an
 104 interpreter for the untyped lambda calculus. While the encoding of type checkers in Datalog
 105 has recently been explored [17], we are the first to support program transformations and
 106 interpreters for Turing-complete languages in Datalog without relying on an embedded
 107 functional programming language. Second, we implemented textbook reaching definitions
 108 and interval analyses. Both analyses are flow-sensitive and compute a fixpoint over the
 109 control-flow graph. Last, we implement clone detection of Java bytecode which is represented
 110 as Soufflé facts. We generate abstract syntax trees by querying Soufflé relations. We then
 111 use the abstract syntax trees to determine if two methods are alpha-equivalent in respect to

112 their identifiers and labels. Our case studies show that functional IncA is expressive and easy
 113 to use. Early performance measurements indicate that reusing established Datalog solvers
 114 yields more efficient execution times.

115 Practically speaking, we consider functional IncA to be a stepping stone for the compilation
 116 of other languages to Datalog. On one hand, our encoding paves the road for transferring
 117 years of research on functional programming languages to Datalog. For example, we show in
 118 this paper how standard defunctionalization [18] can be used to add first-class functions and
 119 first-class relations to functional IncA. Defunctionalization translates first-class functions to
 120 first-order functions and algebraic data types, which we can then compile to Datalog. On
 121 the other hand, we believe that our methodology for supporting user-defined functions and
 122 user-defined data types can be used to compile domain-specific languages to Datalog. We
 123 leave this avenue of research for future work.

124 In summary, we present the following contributions:

- 125 ■ We identify 5 principles that are necessary for the semantics-preserving translation of
 126 first-order functions to Datalog. We define the translation formally and adapt a demand
 127 transformation. This constitutes the first version of functional IncA (Section 3).
- 128 ■ We show how to compile user-defined algebraic data types to Datalog and extend functional
 129 IncA accordingly (Section 4).
- 130 ■ We add sets and set operations to functional IncA, extend the translation, and show
 131 how standard defunctionalization can be used to add first-class functions and first-class
 132 relations (Section 6).
- 133 ■ We demonstrate the expressiveness and ease of use of functional IncA by implementing a
 134 type checker, program transformation, and interpreter for the lambda calculus (Section 5),
 135 data-flow analyses for the While language (Subsection 7.1), and clone detection of Java
 136 bytecode (Subsection 7.2).
- 137 ■ We provide two backends for functional IncA, one targeting the incremental Datalog solver
 138 used by IncA, the other targeting the non-incremental Datalog solver Soufflé (Section 8).

139 2 Datalog Frontends: State of the Art

140 We are by far not the first to recognize the shortcomings of Datalog’s frontend. Two opposing
 141 approaches have been explored in prior work to improve the expressiveness and/or usability
 142 of Datalog. We call these approaches *backend-first* and *frontend-first* and discuss them below.

143 **Backend-first approach.** The backend-first approach uses existing Datalog solvers as a
 144 starting point and extends them with new language features. Usually, extensions considered
 145 in the backend-first approach aim to increase the expressivity of Datalog, but sometimes
 146 also focus on usability. The backend-first approach has a long tradition in Datalog solvers
 147 and some features have become standard nowadays. For example, Datalog solvers support
 148 stratified negation and arithmetic operations, even though neither is part of core Datalog [15].

149 Modern Datalog solvers provide a range of different extensions that their users can choose
 150 from. For example, Soufflé [19] provides records, algebraic data types, and user-defined
 151 functions; Viatra Query [29], the Datalog solver used by IncA, supports user-defined data
 152 types and recursive aggregation over user-defined functions [22, 23]. While all of these
 153 features improve the frontend and make Datalog programming easier, the core language
 154 design remains the same: Horn clauses.

155 Horn clauses ($a_0 :- a_1, \dots, a_n$) encode implications ($a_1 \wedge \dots \wedge a_n \rightarrow a_0$). We argue Horn
 156 clauses are inadequate as a programming language, since they inhibit structured programming
 157 and enforce a flat structure. For example, a nested function call `res = f(g(h(x)))` becomes:

```

158 R(x, res) :- h(x, y), g(y, z), f(z, res)
159

```

161 That is, we must flatten the call chain. Or consider an expression that contains nested condi-
 162 tionals (if (b1) x1 else x2) + (if (b2) x3 else x4), which becomes 4 separate Horn clauses:

```

163 R(b1, b2, x1, x2, x3, x4, res) :- b1, b2, res = x1 + x3.
164 R(b1, b2, x1, x2, x3, x4, res) :- b1, !b2, res = x1 + x4.
165 R(b1, b2, x1, x2, x3, x4, res) :- !b1, b2, res = x2 + x3.
166 R(b1, b2, x1, x2, x3, x4, res) :- !b1, !b2, res = x2 + x4.
167

```

169 These encodings are cumbersome to work with; they make programming and maintenance
 170 unnecessarily difficult. We would much rather use functional programming as a frontend.

171 While the backend-first approach does not fundamentally improve Datalog’s frontend, it
 172 has one decisive advantage: It leverages existing solvers. These solvers are often the result of
 173 years of research and engineering. They automatically optimize Datalog programs, employ
 174 highly optimized data structures and algorithms, support profiling and debugging, provide
 175 incremental execution, and more. When designing new Datalog frontends, we should aim to
 176 reuse these systems. However, the state of the art moves in another direction.

177 **Frontend-first approach.** Quite a few recent research projects try to improve the frontend
 178 of Datalog by designing new DSLs to be used in its stead. We call these approaches frontend-
 179 first because the newly designed frontend is their starting point. In particular, frontend-first
 180 approaches do not build on top of an existing solver but develop a new solver specific to the
 181 newly designed frontend. This allows for great flexibility in the frontend’s design.

182 For example, Flix [14] provides a Datalog frontend extended with lattices and monotonic
 183 functions. Flix embeds its Datalog frontend into a functional programming language, where
 184 constraints are first-class and can be generated at run time [13]. Formulog [9] provides a
 185 Datalog frontend extended with a data type for constructing SMT formulas and a constraint
 186 for solving them. While Formulog constraints are not first-class, the Datalog frontend is also
 187 embedded into a functional programming language. In both Flix and Formulog, the Datalog
 188 constraints can invoke functional code to assert a property or to construct new terms. In
 189 Formulog, functional code can also recursively query Datalog relations. While both systems
 190 present interesting designs, they also both implement their own Datalog solvers and do not
 191 benefit from prior engineering efforts.

192 Datafun [6] proposes a more drastic redesign for Datalog, namely as a higher-order
 193 functional programming language with fixpoint semantics. Datafun functions can accept and
 194 produce relations and the language supports the aggregation over lattices. As such, we believe
 195 Datafun’s frontend is a well-suited replacement for Datalog. However, there are two limiting
 196 factors, First, in contrast to other modern implementations of Datalog, Datafun programs
 197 are constructor-free and enforce termination. While this equips Datafun with a nicer theory,
 198 it is a practical limitation, although one that could be easily eliminated. Second, like Flix
 199 and Formulog above, Datafun provides its own Datalog solver and existing optimizations
 200 and advances in Datalog engines have to be retrofitted to Datafun. For example, semi-naïve
 201 evaluation had to be adapted for Datafun [5], even though it has been the standard bottom-up
 202 evaluation model for a long time [27].

203 **Our approach: Frontend compilation.** We would like to achieve the best of both prior
 204 approaches: Build on top of existing Datalog solvers as in the backend-first approach, but be
 205 free to design functional and domain-specific frontends as in the frontend-first approach. The
 206 solution to this problem is compilation: By compiling the frontend language to Datalog, we
 207 can use existing solvers to run programs. This way, Datalog really becomes the intermediate

208 representation (IR) of a compiler framework, where different Datalog frontends all generate
 209 the same Datalog IR. This architecture is well-known from existing compiler frameworks
 210 such as LLVM; we propose to adopt it for Datalog.

211 Although frontend compilation may seem like the obvious solution, it is difficult to
 212 implement. The problem is that Datalog imposes severe restrictions on programs, so that
 213 bottom-up evaluation is well-defined and terminates. When generating Datalog code, we
 214 must adhere to these restrictions. In the remainder of this paper, we show how a first-order
 215 functional language (Section 3) with algebraic data types (Section 4), and sets (Section 6)
 216 can be compiled to Datalog. In doing so, we will solve key challenges regarding user-defined
 217 functions and user-defined data types that can be transferred to other frontends.

218 3 Compiling First-Order Functions to Datalog

219 We want to provide a functional-programming frontend for Datalog. In this section, we
 220 tackle the first step in this direction: Compiling user-defined first-order functions to Datalog.
 221 While we already outlined why this is challenging in the introduction, here we revisit the
 222 problem with a more involved example before presenting our solution.

223 3.1 Compilation by example

224 In this paper and in our implementation, we use a simple functional frontend language that
 225 features first-order function definitions, let bindings, conditionals, and arithmetic operations.
 226 We also support algebraic data types, set operations, and first-class functions, which we will
 227 explain later. Consider the following recursive factorial function in functional InCA:

```
228 def fact(n: Int): Int = if (n == 0) 1 else n * fact(n - 1)
```

231 We aim to write functions like this and compile them to Datalog, so that we can use them as
 232 part of larger Datalog programs. A simple strategy gets us close to the desired result:

233 **Principle 1: Functions as relations.** It is well-known that functions $f : (T_1, \dots, T_n) \rightarrow T$
 234 can be encoded as relations $f : (T_1, \dots, T_n, T)$. We use this encoding of functions.

235 **Principle 2: Control-flow paths as rules.** For each path from function entry to function
 236 exit, we generate a rule that describes how inputs translates to outputs. Since control-flow
 237 paths are mutually exclusive in deterministic languages, so are the rules we generate.

238 When we apply this strategy to our factorial function, we obtain a relation `fact`: (Int, Int) .
 239 Since the `fact` function has two exits, we derive two rules that collect all conditions and
 240 computations along the path from entry to exit. In doing so, we introduce auxiliary variables
 241 for intermediate results as needed.

```
242 fact(n, out) :- n = 0, out = 1.  
243 fact(n, out) :- n != 0, fact(n-1, out'), out = n * out'.
```

246 Unfortunately, like in the introduction, the Datalog rules violate *range-restrictedness*. A
 247 rule is range-restricted if every variable that occurs in the head of the rule is bound in the
 248 body of the rule. Range-restrictedness is an important property for Datalog programs and
 249 a prerequisite for bottom-up evaluation. Datalog engines like Soufflé [20] apply bottom-
 250 up evaluation to exhaustively enumerate all derivable tuples. Usually, this is an efficient
 251 evaluation strategy, but it diverges for rules that are not range-restricted. In our example,
 252 the second rule is not range-restricted because `n` is not bound in the body, hence `n` could be
 253 any integer term. It follows that the `fact` relation contains infinitely many tuples. Therefore,
 254 Soufflé will reject the Datalog code we generated for the `fact` function.

255 It is hardly surprising that functions over (virtually) infinite domains describe (virtually)
 256 infinite relations. So is this approach doomed? To move forward, we make an important
 257 observation: Even though a function may be defined over an infinite domain, *any terminating*
 258 *application of that function will only see finitely many inputs*. If we can restrict a function's
 259 relation to these inputs, the entire relation turns finite and each rule becomes range-restricted.

260 To determine the relevant inputs of a function, we must consider how the function is
 261 used and what inputs it is applied to. For our factorial example, consider a main call `fact(5)`,
 262 which stipulates that `n = 5` is a relevant input of the `fact` relation. But since `fact` is recursive,
 263 we must also track which relevant inputs are induced by `n = 5`. If we collect all relevant
 264 inputs in `fact_input = {5,4,3,2,1,0}`, we can use this relation to guard the bodies of `fact`:

```
265 run_fact(out) :- fact(5,out).
266 fact(n, out) :- fact_input(n), n = 0, out = 1.
267 fact(n, out) :- fact_input(n), n != 0, fact(n-1, out'), out = n * out'.
268
```

270 Note how all rules are range-restricted now. Input variables are range-restricted by the query
 271 of the input relation; output variables are range-restricted because they are functionally
 272 dependent on the input variables. Thus, `fact` is finite when `fact_input` is finite.

273 **Principle 3: Input relations as guards.** For each function, collect all relevant inputs in an
 274 input relation and use the input relation as a guard for the function's relation.

275 Relevant inputs stem from external calls of the function or from recursive calls. Therefore, it
 276 is not easy to collect all relevant inputs in a relation. Fortunately, we can apply an existing
 277 algorithm that is well-known in the Datalog community: the magic-set transformation [8]. The
 278 magic-set transformation was developed to optimize the bottom-up evaluation of terminating
 279 Datalog programs. The key idea of the magic-set transformation is to only derive those tuples
 280 bottom-up that would also be derived by top-down evaluation, where the relevant inputs
 281 are known. To this end, the magic-set transformation generates Datalog rules for auxiliary
 282 relations that prescribe which inputs are relevant. Note that we say "inputs" here because
 283 the relations we care about correspond to functions; in general, the magic-set transformation
 284 collects terms that are known at the call-site during run time. Since function inputs are
 285 always known at the call-site during run time, the magic-set transformation will at least
 286 collect all relevant function inputs. Technically, we apply a more efficient variation of the
 287 magic-set transformation called the *demand transformation* [25] and we use that name in
 288 the remainder of the paper.

289 **Principle 4: Demand transformation yields input relations.** The demand transformation
 290 identifies all relevant inputs for each function in the program. Since all function call-sites
 291 must be known, our compilation strategy is not modular but requires the whole program.

292 For our example, the demand transformation will generate the following input relation:

```
293 fact_input(5).
294 fact_input(n-1) :- fact_input(n), n != 0.
295
296
```

297 We obtain one rule for each call of `fact`. The first rule collects the input of the main
 298 invocation `fact(5)`. The second rule collects the input of the recursive invocation and contains
 299 all constraints leading up to the call. Together, these two rules describe the required relation
 300 `fact_input = {5,4,3,2,1,0}`. Since `fact_input` is finite, `fact` is finite and contains the following
 301 tuples: `fact = {(5,120), (4,24), (3,6), (2,2), (1,1), (0,1)}`.

302 So far, all function inputs were statically known. But we can easily extend our compilation
 303 strategy to support user-provided inputs. To this end, functional IncA allows the declaration
 304 of main functions:

```
305 @main def run_fact(n: Int): Int = fact(n)
306
```

(Functional programs)	$p ::= \overline{F}$
(functions)	$F ::= [\text{@main}] \text{def } f(\overline{x:T}): T = e$
(expressions)	$e ::= v \mid x \mid \text{let } x = e \text{ in } e \mid \text{if } (e) e \text{ else } e \mid f(\overline{e}) \mid \varphi(\overline{e})$
(values)	$v ::= \text{base}$
(types)	$T ::= \text{Base}$

■ **Figure 3** Functional IncA with first-order functions, base values, and base functions φ .

(Datalog programs)	$D ::= \overline{r}$
(rules)	$r ::= \mathbf{R}(\overline{t}) :- \overline{a}.$
(atoms)	$a ::= t = t \mid \mathbf{R}(\overline{t})$
(terms)	$t ::= v \mid x \mid \varphi(\overline{t})$
(values)	$v ::= \text{base}$

■ **Figure 4** An intermediate representation for Datalog with base values and base functions.

308 The demand transformation will correctly propagate the input of `run_fact` to `fact`:

```
309 fact_input(n) :- run_fact_input(n).
310 fact_input(n-1) :- fact_input(n), n != 0.
311
```

313 But what is the input of `run_fact`? The input of `run_fact` is dynamic and must be provided
 314 by the user of the program. In Datalog, such data lives in the so-called extensional database,
 315 which is filled by the user prior to Datalog execution. We modify the demand transformation
 316 to generate a query of the extensional database for main functions.

317 **Principle 5: Main input in extensional database.** For each main function, we add a rule to
 318 the input relation that retrieves dynamic inputs from the extensional database.

319 For our example, we obtain the following input relation for `run_fact`:

```
320 run_fact_input(n) :- ext_run_fact_input(n).
321
```

323 The user can provide any number of inputs to `run_fact` as part of the extensional database.
 324 The Datalog engine will propagate those inputs to `run_fact_input` and fill all relations.

325 Note that our encoding retains crucial Datalog behavior, such as memoization and reuse.
 326 For example, consider we want to run `fact` on multiple inputs 5, 7, and 9, all of which we put
 327 into the extensional database. How many tuples will relation `fact` contain? Since queries
 328 of `fact` will retrieve existing tuples when possible, the three `fact` computations will share
 329 all intermediate results and `fact` will only contain 10 tuples (the largest input plus one). A
 330 similar effect can be observed for functions like Fibonacci, where recursive calls can share
 331 results. All of this is transparent to the user.

322 3.2 Translating functional programs to Datalog, technically

333 We now implement Principles 1 and 2 from the previous subsection, that is, we translate
 334 functional programs to Datalog. In the subsequent subsection, we will explain and apply the
 335 demand transformation to implement the remaining principles.

336 Figure 3 defines the syntax of functional IncA. The language consists of first-order
 337 functions, let bindings, conditionals, and function calls. We distinguish calls to user-defined
 338 functions f from calls to base functions φ . Our compilation target is an intermediate
 339 representation (IR) of Datalog extended with base values and base functions as shown in
 340 Figure 4. This Datalog IR is compatible with many existing Datalog solvers, which support
 341 different kind of base functions. Note that we excluded negation from the Datalog IR because
 342 our translation does not require it.

343 We first translate expressions to Datalog. While an expression is structured and eventually
 344 computes a value, Datalog only provides flat terms. Thus, a nested expression $f(g(x))$ must be

$$\begin{aligned}
\llbracket \cdot \rrbracket &: e \rightarrow \mathcal{P}(t \times \mathcal{P}(a)) \\
\llbracket v \rrbracket &= \{(v, \emptyset)\} \\
\llbracket x \rrbracket &= \{(x, \emptyset)\} \\
\llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket &= \{(t_2, \{x = t_1\} \cup a_1 \cup a_2) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, (t_2, a_2) \in \llbracket e_2 \rrbracket\} \\
\llbracket \text{if } (e_1) e_2 \text{ else } e_3 \rrbracket &= \{(t_2, \{t_1 = \text{true}\} \cup a_1 \cup a_2) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, (t_2, a_2) \in \llbracket e_2 \rrbracket\} \\
&\quad \cup \{(t_3, \{t_1 = \text{false}\} \cup a_1 \cup a_3) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, (t_3, a_3) \in \llbracket e_3 \rrbracket\} \\
\llbracket f(e_1, \dots, e_n) \rrbracket &= \{(y, \{f(t_1, \dots, t_n, y)\} \cup a_1 \cup \dots \cup a_n) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, \dots, (t_n, a_n) \in \llbracket e_n \rrbracket\} \\
&\quad \text{where } y \text{ is fresh} \\
\llbracket \varphi(e_1, \dots, e_n) \rrbracket &= \{(\varphi(t_1, \dots, t_n), a_1 \cup \dots \cup a_n) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, \dots, (t_n, a_n) \in \llbracket e_n \rrbracket\}
\end{aligned}$$

■ **Figure 5** Compiling expressions yields a set of alternative terms, each guarded by constraints.

$$\begin{aligned}
\llbracket \text{def } f(\overline{x:T}) : T' = e \rrbracket_{fun} &= \{\mathbf{f}(\overline{x}, y) :- a, y = t. \mid (t, a) \in \llbracket e \rrbracket\} \quad \text{where } y \text{ is fresh} \\
\llbracket \overline{F} \rrbracket_{prog} &= \bigcup_{f \in \overline{F}} \llbracket f \rrbracket_{fun}
\end{aligned}$$

■ **Figure 6** Compiling functions to Datalog rules.

345 compiled to a flat term that is guarded by constraints $(y_2, \{g(x, y_1), f(y_1, y_2)\})$. Since condi-
346 tional expressions (if $(b) f(x) \text{ else } g(x)$) yield alternative values depending on b , compilation in
347 general yields a set of alternative terms $\{(y_1, \{b = \text{true}, f(x, y_1)\}), (y_2, \{b = \text{false}, g(x, y_2)\})\}$.
348 This corresponds to Principle 2 from the previous subsection.

349 Figure 5 defines the translation of expressions as a compositional function $\llbracket \cdot \rrbracket$. Values v
350 and variables x directly translate to Datalog values and variables. Let bindings yield the
351 body's result under a constraint that binds the let-bound variable. Conditionals compile to
352 two alternative sets of terms: If the condition is **true**, the resulting terms are taken from the
353 *then*-branch, otherwise they are taken from the *else*-branch. Calls to user-defined functions
354 f translate to queries of a relation of the same name f , which has the function's result as
355 an additional column in accordance with Principle 1. In contrast, calls to base functions φ
356 translate to a call of the same function, but passing Datalog terms as arguments.

357 We use the translation of expressions $\llbracket \cdot \rrbracket$ to compile function definitions $\llbracket \cdot \rrbracket_{fun}$ and
358 programs $\llbracket \cdot \rrbracket_{prog}$ as shown in Figure 6. For a function definition, we compile its body and
359 generate a separate Datalog rule for each alternative term that the body can yield. The
360 constraints a of the term become constraints in the generated rule. To compile a whole
361 program, we simply compile each function and collect the resulting rules.

362 For a concrete example, consider the translation of the Fibonacci function to Datalog:

363 $\llbracket \text{def fib}(n) = \text{if } (n < 2) \ 1 \ \text{else fib}(n-1) + \text{fib}(n-2) \rrbracket = \{ \text{fib}(n, y_3) :- n < 2 = \text{true}, y_3 = 1.,$
364 $\text{fib}(n, y_4) :- n < 2 = \text{false}, \text{fib}(n-1, y_1), \text{fib}(n-2, y_2), y_4 = y_1 + y_2. \}$
365

367 The Fibonacci function compiles to two Datalog rules, one for the base case and one for the
368 recursive case. But is this translation correct?

369 **Translation correctness.** We claim that our translation preserves the semantics of the
370 original functional program. More precisely, we claim that if a function call $f(\overline{v})$ evaluates
371 to w , then the generated Datalog program will also provide w as the only result of the call
372 *under top-down evaluation*. Here we must require top-down evaluation for Datalog, since
373 the generated rules are not necessarily range-restricted yet, which we will fix in the next
374 subsection. Top-down evaluation is possible nonetheless, because it only explores required
375 results and uses known values in doing so. Since the values of function arguments are always

XX:10 Functional Programming with Datalog

376 known during evaluation, top-down evaluation of the generated Datalog closely corresponds
377 to function evaluation. However, we did not formalize top-down evaluation and therefore
378 formulate translation correctness as a conjecture:

379 ► **Conjecture 1** (Translation correctness). *Given a functional program p with a main function*
380 *f such that $f(\bar{v})$ evaluates to w . Then the top-down evaluation of the Datalog atom $f(\bar{v}, x)$*
381 *under $\llbracket p \rrbracket_{prog}$ yields a single substitution $\{x \mapsto w\}$.*

382 A key component for proving this conjecture is to ensure the Datalog constraints behave
383 deterministically, just like the original expression did:

384 ► **Lemma 2** (Deterministic atoms). *Given an expression e such that $\llbracket e \rrbracket = \{(t_1, \bar{a}_1), \dots, (t_n, \bar{a}_n)\}$.*
385 *Both of the following hold:*

- 386 *i. $\llbracket e \rrbracket$ yields at least one result: $\bar{a}_1 \vee \dots \vee \bar{a}_n$*
387 *ii. $\llbracket e \rrbracket$ yields at most one result: $(\bar{a}_i \wedge \bar{a}_j) \rightarrow t_i = t_j$*

388 **Proof.** By structural induction over e . The only interesting case are *if*-expressions, where
389 $(t_1 = \text{true})$ and $(t_1 = \text{false})$ are mutually exclusive. ◀

390 ► **Lemma 3** (Deterministic rules). *Given f such that $\llbracket f \rrbracket_{fun} = \{f(\bar{x}, y_1) :- \bar{a}_1, \dots, f(\bar{x}, y_n) :- \bar{a}_n\}$.*
391 *Both of the following hold:*

- 392 *i. $\llbracket f \rrbracket_{fun}$ yields at least one result: $\bar{a}_1 \vee \dots \vee \bar{a}_n$*
393 *ii. $\llbracket f \rrbracket_{fun}$ yields at most one result: $(\bar{a}_i \wedge \bar{a}_j) \rightarrow y_i = y_j$*

394 **Proof.** Follows from Lemma 2. ◀

395 Note that Conjecture 1 does not make any assertions about non-terminating function calls.
396 Indeed, some diverging functions compile to terminating Datalog programs. For example,
397 `def f(x) = f(x)` compiles to `f(x,y) :- f(x,y)`. While function call `f(1)` diverges, query `f(1, y)`
398 terminates and yields the empty substitution. However, Conjecture 1 ensures terminating
399 function calls translate to terminating Datalog programs under top-down evaluation.

400 3.3 Demand-driven bottom-up evaluation

401 We compile functional programs to Datalog rules that execute well in top-down fashion,
402 but may diverge under bottom-up evaluation. In bottom-up evaluation, Datalog solvers
403 exhaustively enumerate all derivable tuples, starting from known facts. For example, the
404 bottom-up evaluation of the factorial function will start with `fact(0,1)`, from which it can
405 derive `fact(1,1)`, `fact(2,2)`, `fact(3,6)`, `fact(4,24)`, and so on. This enumeration will not
406 terminate, because bottom-up evaluation is unaware of the context in which relation `fact`
407 is being used. Accordingly, we cannot apply any of the efficient Datalog solvers that use
408 bottom-up evaluation, such as Soufflé.

409 The demand transformation by Tekle and Liu rewrites Datalog rules such that bottom-up
410 evaluation becomes demand-driven and only computes tuples that are transitively demanded
411 by the main query [25]. Indeed, bottom-up evaluation of the rewritten Datalog rules computes
412 *exactly the same* tuples as top-down evaluation. Since we already asserted that top-down
413 evaluation computes the correct result for terminating functional programs, the demand
414 transformation allows us to apply bottom-up evaluation, also yielding the correct result.

415 We adopt the demand transformation, which transforms a set of Datalog rules in three
416 steps: compute demand patterns, introduce demand predicates, derive demand rules. In
417 this section, we replace the first step of the demand transformation to take functional IncA
418 into account, adopt the second step unchanged, and extend the third step to account for the
419 inputs of main functions. Later sections will make further changes.

420 **Step 1** We compute demand patterns $\langle g, s \rangle$, where g is the name of a relation and $s \in (b | f)^*$
 421 is a pattern string that indicates how the relation is queried, namely if an argument occurs
 422 bound or free. For functions, demand patterns can be easily computed by finding all function
 423 calls reachable from the main functions. Formally, given a functional program p , the demand
 424 patterns $dp(p)$ of p is the smallest set such that:

- 425 ■ For each main function (`@main def $g(x_1, \dots, x_n) = \dots$`) in p , we have $\langle g, b^n f \rangle \in dp(p)$.
 426 That is, main functions have demand with n bound parameters and one free return value.
- 427 ■ If demand pattern $\langle g, s \rangle \in dp(p)$ and g is defined as (`def $g(\dots) = e$`) in p , we have
 428 $\langle h, b^n f \rangle \in dp(p)$ for each call $h(e_1, \dots, e_n)$ in e .

429 The second and third step of the demand transformation operate on and rewrite the generated
 430 Datalog rules $D = \llbracket p \rrbracket_{prog}$. In particular, we will make no assumptions about the format of
 431 pattern strings s , so that we can later introduce extensions of Step 1 easily.

432 **Step 2** We introduce demand predicates as guards into existing rules to implement Princi-
 433 ple 3 from Subsection 3.1. Formally, we obtain a rewritten Datalog program $guarded(D)$:

- For each $\langle g, s \rangle \in dp(p)$ and each $(g(t_1, \dots, t_m) :- a_1, \dots, a_n.)$ in D , we obtain a rule

$$g(t_1, \dots, t_m) :- \underline{g_input_s}(t_1, \dots, t_m|_s), a_1, \dots, a_n.$$

434 in $guarded(D)$, where $\bar{t}|_s$ selects those t_i that are bound according to pattern string s .
 435 Note that the rules of unreachable functions are dropped and not propagated to $guarded(D)$.

436 **Step 3** In the final step, we must derive those rules that define the input relations $\underline{g_input_s}$
 437 to implement Principle 4 and Principle 5 from Subsection 3.1. Formally, we obtain a rewritten
 438 Datalog program $demanded(D)$ from $guarded(D)$ and the original program p as follows:

- 439 ■ We retain each rule from $guarded(D)$, such that $guarded(D) \subseteq demanded(D)$.
- For each main function (`@main def $g(x_1, \dots, x_n) = \dots$`) in p , we obtain a rule

$$\underline{g_input_s}(x_1, \dots, x_n) :- \underline{ext_g_input_s}(x_1, \dots, x_n).$$

440 in $demanded(D)$, where $\underline{ext_g_input_s}$ is an extensional relation to be filled by the user.
 441 This implements Principle 5.

- For each rule $(g(\dots) :- a_1, \dots, a_n.)$ in $guarded(D)$ and each $a_i = h(t_1, \dots, t_m)$, we obtain

$$\underline{h_input_s}(t_1, \dots, t_m|_s) :- a_1, \dots, a_{i-1}$$

442 to $demanded(D)$, where s is the pattern string of $h(t_1, \dots, t_m)$, indicating which t_i are
 443 bound by the previous constraints a_1, \dots, a_{i-1} already.

444 The demand transformation implements Principles 3 - 5 and ensures that the resulting
 445 Datalog derives the same tuples in bottom-up evaluation as in top-down fashion.

446 **Example** To illustrate, consider again the Fibonacci function with a main call:

```
447 def fib(n) = if (n<2) 1 else fib(n-1) + fib(n-2)
448 @main def run(x: Int): Int = fib(x)
449
450
```

451 This program compiles to the following Datalog rules using the translation from Subsection 3.2:

XX:12 Functional Programming with Datalog

```
452 fib(n, y3) :- n<2 = true, y3 = 1.
453 fib(n, y4) :- n<2 = false, fib(n-1,y1), fib(n-2,y2), y4 = y1+y2.
454 run(x, y5) :- fib(x, y5).
```

457 We now apply our demand transformation. First, we derive demand patterns of the program,
458 which are $\langle \text{run}, bf \rangle$ and $\langle \text{fib}, bf \rangle$. Note that all three calls of `fib` yield the same demand pattern.
459 Second, we insert demand predicates into the rules according to the demand patterns:

```
460 fib(n, y3) :- fib_input_bf(n), n<2 = true, y3 = 1.
461 fib(n, y4) :- fib_input_bf(n), n<2 = false, fib(n-1,y1), fib(n-2,y2), y4 = y1+y2.
462 run(x, y5) :- run_input_bf(x), fib(x, y5).
```

465 Third, to these rules we add the following rules to define the input relations:

```
466 run_input_bf(x) :- ext_run_input_bf(x).
467 fib_input_bf(x) :- run_input_bf(x).
468 fib_input_bf(n-1) :- fib_input_bf(n), n<2 = false.
469 fib_input_bf(n-2) :- fib_input_bf(n), n<2 = false, fib(n-1,y1).
```

472 The first and second rule are due to the main function `run`, which receives its input from
473 the user and propagates it to `fib`. The third and fourth rule are due to the recursive calls
474 of `fib`. Note how we retain all constraints prior to a call. In particular, we retain the first
475 recursive call of `fib` as a constraint for the second recursive call of `fib`, although a smart
476 compiler might eliminate this constraint subsequently. The resulting Datalog program is
477 demand-driven and can be executed by standard bottom-up Datalog solvers.

478 **Correctness** The demand transformation yields a Datalog program that derives the exact
479 same tuples as a top-down evaluation [25]. As of Conjecture 1, top-down evaluation yields
480 the correct tuples. Hence, so does bottom-up evaluation of the demand-driven Datalog rules:

481 ► **Corollary 4** (Bottom-up translation correctness). *Given a functional program p with a main*
482 *function f such that $f(\bar{v})$ evaluates to w . Then the bottom-up evaluation of the Datalog*
483 *program $\text{demanded}(\llbracket p \rrbracket_{prog})$ yields a database in which the query $f(\bar{v}, x)$ has a single match*
484 *$\{f(\bar{v}, w)\}$.*

4 Compiling Algebraic Data Types to Datalog

486 The functional IncA we presented in the previous section supports user-defined functions
487 ranging over base types. In this section, we explore how to extend functional IncA to allow
488 user-defined data types. In particular, we want to faithfully compile recursive functions over
489 algebraic data types to Datalog rules that existing bottom-up Datalog solvers can execute.

4.1 Compiling user-defined data types by example

491 We extend functional IncA to allow recursive definitions of user-defined algebraic data types,
492 constructor calls, and pattern matching. As a simple example, consider the Peano numbers:

```
493 data Nat = Zero() | Succ(Nat)
494 def plus(m: Nat, n: Nat): Nat = m match {
495   case Zero() => n
496   case Succ(pred) => Succ(plus(pred, n)) }
497 @main def twice(n: Nat): Nat = plus(n, n)
```

500 We generate three kind of relations for an algebraic data type:

- 501 ■ **Constructor relations** represent the constructor functions of algebraic data types. We
- 502 translate constructor calls in the program to queries of constructor relations, similar to
- 503 how we translated regular function calls. In doing so, it is crucial we ensure only finitely
- 504 many values are constructed during bottom-up evaluation of the resulting Datalog code.
- 505 ■ **Selector relations** map a constructed value to its constituents. We use selector relations
- 506 to implement pattern matching. Importantly, queries of selector relations may never lead
- 507 to the construction of new values.
- 508 ■ **Instance relations** enumerate all constructed instances of a data type. They will become
- 509 useful when we introduce relational programming in Section 6.

510 To construct user-defined data at run time, we extend the Datalog IR with a built-in
 511 constructor `#constr` for each constructor `constr`. For example, `#Succ(#Succ(#Zero()))` encodes
 512 two as a Peano number. In practice, there are different ways a Datalog solver can support
 513 such built-in constructors. For example, we can define a generic built-in function that creates
 514 a new value given the constructor's name and arguments. We have used this approach in
 515 our implementation using IncA, but this would work in any Datalog solver that supports
 516 user-defined built-in functions, including Soufflé, Flix, and Formulog. Alternatively, if a
 517 Datalog solver natively supports algebraic data types, we can use their constructors directly
 518 or encode them using a number representation. For example, Soufflé supports algebraic data
 519 types (but not recursive functions over them) and we can generate a Soufflé data type and
 520 use its constructors. This is to say that adding built-in constructors to the Datalog IR does
 521 not limit the applicability of our approach in practice. Flix, Formulog, IncA and Soufflé have
 522 support for algebraic data. However, they do not support enumerating all instances of a
 523 specific algebraic data type like functional IncA. We will see how to enumerate all instances
 524 of an algebraic data type by utilizing instance relations in Section 6.

525 For the Peano numbers, we derive the following Datalog rules initially:

```

526 // constructor relations      // selector relations      // instance relation
Zero(n) :- n = #Zero().      un_Zero(n) :- Zero(n).      Nat(n) :- Zero(n).
Succ(p, n) :- n = #Succ(p).  un_Succ(n, p) :- Succ(p, n). Nat(n) :- Succ(_, n).

```

527 Note that the rule of the `Succ` constructor relation is not range-restricted and consequently
 528 cannot be computed bottom-up. However, the rules of the selector and instance relations
 529 merely query the constructor relations. Hence, if we can ensure the constructor relations
 530 remain finite, all three kind of relations will be finite.

531 Like in the previous section, we seek to apply the demand transformation in order to track
 532 the demand of constructor relations. However, we need to adapt the demand transformation
 533 to account for our encoding of algebraic data types. Specifically, the constructor queries
 534 within the selector and instance relations must be ignored, since they do not actually indicate
 535 additional demand. Moreover, selector and instance relations do not require any rewriting
 536 themselves, because they merely query constructor relations to enumerate constructor tuples.

537 Figure 7 shows the compilation result after demand transformation for the `pplus` function on
 538 Peano numbers from above. Relation `Zero` has no demand relation because its demand pattern
 539 $\langle \text{Zero}, f \rangle$ does not specify bound inputs. Relation `Succ` has a demand relation `Succ_input_bf`
 540 that tracks the invocation of `Succ` in the recursive case of `pplus`. Importantly, there is no
 541 demand on `Succ` from the selector or instance relations, as our adaption of the demand
 542 transformation will ensure. Relation `pplus` shows how we compile pattern matching: Each
 543 case becomes an alternative rule that queries the selector. This is sufficient since we assume
 544 pattern matches are complete and overlap-free, so that their order does not matter.

545 Since `twice` is a main function, its demand relation queries an extensional input relation as
 546 described in the previous section. This way, users can for example request `twice(Succ(Zero()))`.

```

Zero(n) :- n = #Zero(). // no demand relation since there are no bound inputs
Succ(p, n) :- Succ_input_bf(p), n = #Succ(p).
Succ_input_bf(y4) :- plus_input_bbf(m, n), un_Succ(m, pred), plus(pred, n, y4).

// selector and instance relations un_Zero, un_Succ, and Nat as above
plus(m, n, out) :- plus_input_bbf(m, n), un_Zero(m), out = n.
plus(m, n, out) :- plus_input_bbf(m, n), un_Succ(m, pred),
                    plus(pred, n, y4), Succ(y4, y5), out=y5.
plus_input_bbf(n, n) :- twice_input_bf(n).
plus_input_bbf(pred, n) :- plus_input_bbf(m, n), un_Succ(m, pred).

twice(n, out) :- twice_input_bf(n), plus(n, n, out).
twice_input_bf(n) :- ext_twice_input_bf(n).

Zero(n) :- ext_Zero(n).
Succ(p, n) :- ext_Succ(p, n).

```

■ **Figure 7** Compilation result for the `plus` and `twice` functions on Peano numbers.

(Functional programs)	$prog ::= \overline{F}, \overline{d}$
(data definitions)	$d ::= \text{data } N = \overline{c(T, \dots, T)}$
(expressions)	$e ::= \dots \mid c(\overline{e}) \mid e \text{ match } \{\text{case } c(x, \dots, x) => e\}$
(types)	$T ::= \dots \mid N$

■ **Figure 8** Extending the frontend syntax with algebraic data types.

547 But how can our Datalog program deconstruct the user-provided data? Recall that selector
548 relations simply query constructor relations. Thus, we must include the user-provided
549 algebraic data in our constructor relations. To this end, we require users to insert algebraic
550 data in extensional constructor relations. We then generate one additional rule for each
551 constructor that queries the corresponding extensional constructor relation, as shown at the
552 end of Figure 7. We need to provide the contents of extensional constructor relations in the
553 form of tuples consistent with the format supported by the targeted Datalog dialect. In the
554 case of Soufflé, we insert tuples containing algebraic data and literal values of the Soufflé
555 language in the extensional constructor relations.

556 4.2 Extending functional IncA with algebraic data types

557 Based on the observations from the previous subsection, we add algebraic data types to
558 functional IncA. We then extend the translation from functional code to Datalog code and
559 the demand transformation accordingly.

560 Figure 8 extends the abstract syntax of functional IncA with algebraic data types. For
561 pattern matching we assume that patterns are complete and overlap-free. We do not change
562 the syntax of Datalog since we model constructors as built-in functions φ .

563 We extend the translation of Subsection 3.2 from functional code to Datalog code to
564 handle algebraic data types as shown in Figure 9. We add a new translation function $\llbracket \cdot \rrbracket_{data}$
565 for data types and use that when compiling programs in $\llbracket \cdot \rrbracket_{prog}$. The translation of functions
566 $\llbracket \cdot \rrbracket_{fun}$ remains the same, but it uses an extended translation for expressions $\llbracket \cdot \rrbracket$ that handles
567 the new expressions: constructor calls and pattern matching. The translation of constructor
568 calls is identical to the translation of regular function calls, except the generated code queries
569 a constructor relation. Pattern matching yields alternative rules for each case, and each case
570 queries the selector relation `un_c` to test if the term matches the pattern. The translation of
571 data types $\llbracket \cdot \rrbracket_{data}$ generates rules as described in the previous subsection: rules that invoke
572 the built-in constructor functions, rules that query the extensional constructor relations,

$$\llbracket \bar{F}, \bar{d} \rrbracket_{prog} = \bigcup_{f \in \bar{F}} \llbracket f \rrbracket_{fun} \cup \bigcup_{d \in \bar{d}} \llbracket d \rrbracket_{data}$$

$$\llbracket c(e_1, \dots, e_n) \rrbracket = \{(y, \{c(t_1, \dots, t_n, y)\} \cup a_1 \cup \dots \cup a_n) \mid (t_1, a_1) \in \llbracket e_1 \rrbracket, \dots, (t_n, a_n) \in \llbracket e_n \rrbracket\}$$

where y is fresh

$$\llbracket e \text{ match } \{\bar{cs}\} \rrbracket = \bigcup_{(\text{case } c(\bar{x}) \Rightarrow e') \in \bar{cs}} \{(t', \{\text{un}_c(t, \bar{x})\} \cup a \cup a') \mid (t, a) \in \llbracket e \rrbracket, (t', a') \in \llbracket e' \rrbracket\}$$

$$\begin{aligned} \llbracket \text{data } N = \bar{C} \rrbracket_{data} &= \{c(x_1, \dots, x_n, y) :- y = \#c(x_1, \dots, x_n). \mid c(T_1, \dots, T_n) \in \bar{C}\} \\ &\cup \{c(x_1, \dots, x_n, y) :- y = \text{ext}_c(x_1, \dots, x_n, y). \mid c(T_1, \dots, T_n) \in \bar{C}\} \\ &\cup \{\text{un}_c(y, x_1, \dots, x_n) :- c(x_1, \dots, x_n, y). \mid c(T_1, \dots, T_n) \in \bar{C}\} \\ &\cup \{N(y) :- c(x_1, \dots, x_n, y). \mid c(T_1, \dots, T_n) \in \bar{C}\} \end{aligned}$$

■ **Figure 9** Translating algebraic data types to Datalog.

573 rules for the selector relations, and rules for the instance relations.

574 Next, we extend the demand transformation from Subsection 3.3 to consider constructors:

- 575 ■ In Step 1, when considering reachable subexpressions $h(e_1, \dots, e_n)$, we also generate a
- 576 demand pattern $\langle h, b \dots bf \rangle$ when g is a constructor.
- 577 ■ In Step 2, note that selector and instance relations are never demanded, since we ignored
- 578 them in Step 1. Hence, we propagate their rules unchanged to *guarded*(D).
- 579 ■ In the last case of Step 3, we ignore atoms $a_i = h(t_1, \dots, t_m)$ that occur in the rules of
- 580 selector or instance relations. These atoms always query a constructor relation and we
- 581 do not want to treat these queries as demand.

582 With these modifications, the demand transformation will correctly track the demand

583 of constructors while ignoring selectors and instance relations. Together, the extended

584 translation and the demand transformation constitute a compiler for functional IncA with

585 algebraic data types. Since all rules of the generated Datalog code are range-restricted, we

586 can run the code with off-the-shelf bottom-up Datalog solvers.

587 **5 Case study: Type Checking, Type Erasure, and Interpretation**

588 Functional IncA supports user-defined functions and data types. In this section, we demon-

589 strate that these features allow us to express interesting computations in Datalog. In

590 particular, we implement a type checker, type erasure, and an interpreter for a lambda

591 calculus with numbers as illustrated in Figure 10. These functions compile to complex

592 Datalog code that could not practically be written by hand.

593 Figure 10 shows an excerpt of the relevant data types and functions, all of which are

594 completely standard. In particular, we describe the expressions of the simply typed lambda

595 calculus Exp and the untyped lambda calculus uExp as algebraic data types. We define a

596 type checker `typeOf` as a function in functional IncA, but only show the `App` case here. Our

597 implementation supports parametric polymorphism by applying monomorphization before

598 translating to Datalog. Since the `App` case has five alternative control-flow paths, this case

599 alone compiles into five Datalog rules for `typeOf`. For example, consider the rule generated

600 for the path that yields `Just(ty2)`:

```
601 typeOf(ctx, exp, out0) :-
602   typeOf_input_bbf(ctx, exp), un_App(exp, fun, arg), typeOf(ctx, fun, o1),
603   un_JustType(o1, funty), un_TFun(funty, ty1, ty2), typeOf(ctx, arg, o2),
604   un_JustType(o2, argty), eqType(argty, ty1, o3), o3 == true, JustType(ty2, out0).
```

```

data Exp = Num(Int) | Lam(String, Type, Exp) | App(Exp, Exp) | Var(String)
data Type = TInt() | TFun(Type, Type)
data UExp = UNum(Int) | ULam(String, Exp) | UApp(Exp, Exp) | UVar(String)
def typeOf(ctx: Ctx, exp: Exp): Maybe[Type] = exp match {
  case App(fun, arg) => typeOf(ctx, fun) match {
    case Just(TFun(ty1, ty2)) => typeOf(ctx, arg) match {
      case Just(argty) => if (eqType(argty, ty1)) Just(ty2) else Nothing()
    ... }
  }
}
def erase(exp: Exp): UExp = exp match {
  case Num(v) => UNum(v)
  case Lam(n, ty, b) => ULam(n, erase(b))
  case App(fun, arg) => UApp(erase(fun), erase(arg))
  case Var(n) => UVar(n) }
def interp(env: Env, exp: UExp): Maybe[Val] = exp match {
  case UApp(fun, arg) => interp(env, fun) match {
    case Just(VClosure(param, prog, fenv)) => interp(env, arg) match {
      case Just(argv) => interp(BindEnv(param, argv, fenv), body)
    ... }
  }
}
@main def run(exp: Exp): Maybe[Val] = typeOf(EmptyCtx(), exp) match {
  case Just(ty) => interp(EmptyEnv(), erase(exp))
  case Nothing() => Nothing() }

```

■ **Figure 10** A type checker, type erasure, and interpreter for a lambda calculus with numbers.

607 This Datalog rule consists of 10 atoms, where the selector predicates ensure that the correct
608 control-flow path has been chosen. Overall, the `typeOf` function consists of 24 lines of code,
609 but compiles to 114 lines of complex Datalog code with mutually dependent relations `typeOf`
610 and `typeOf_input`. These numbers represent the Datalog program after applying optimizations.
611 In contrast to program optimizations of functional and imperative programs, our Datalog
612 optimizations reduce the number of rules and atoms instead of increasing them.

613 Next, we define type erasure as a transformation from `Exp` to `UExp`. Although function
614 `erase` is completely standard, this is the first program transformation implemented in Datalog
615 to the best of our knowledge. While `erase` is guaranteed to terminate, we can also define
616 functions whose termination is undecidable. Specifically, we implement a standard interpreter
617 `interp` for the untyped lambda calculus, which is a Turing-complete language. Indeed, the
618 Datalog program is only guaranteed to terminate when the original interpreter terminates.

619 Overall, the type checker, type erasure, and interpreter comprise 8 algebraic data types
620 and 7 functions. We compile this code to 65 relations defined by 154 rules that contain
621 484 atoms in total. These numbers are measured after optimization, where we eliminate
622 aliases and propagate constants.

623 Although implementing an interpreter in Datalog may seem to be of little use, this and
624 similar challenges occur during program analysis regularly. For example, Pacak et al. recently
625 have shown how to compile typing rules to Datalog to derive incremental type checkers
626 systematically [17]. They also mention that it is necessary to translate the dynamic semantics
627 of a language to Datalog in order to support the incremental type checking of a dependently
628 typed programming language. Similarly, data-flow analyses often need to abstractly interpret
629 programs, for example, to determine the bounds of numeric variables or the value of a
630 Boolean condition. Functional IncA can also support such data-flow analyses, but we must
631 be able to express control-flow graphs and other relations.

632 **6** Mixing Functions and Relations

633 The previous sections showed how we can use functional programming as a frontend for
634 Datalog. However, in doing so, we have also lost a key feature of Datalog: relations. Indeed,

```

data Exp = ...
data Stm = Assign(String, Exp) | Sequence(Stm, Stm) | If(Exp, Stm, Stm) | While(Exp, Stm)
def init(stm: Stm): Stm = ... // a regular function that finds the statement's entry
def final(stm: Stm): Set[Stm] = stm match { // finds all of the statement's exits
  case Assign(x, a) => {stm}
  case Sequence(s1, s2) => final(s2)
  case If(b, s1, s2) => final(s1) ++ final(s2)
  case While(b, s) => {stm} }
// flow as seen in Figure 1 (Introduction)

```

■ **Figure 11** Computing the control-flow graph as a set of tuples in our extended Datalog frontend.

functional InCA makes it difficult to encode non-functional relations, such as the edges of a graph. In the present section, we show how we can elegantly extend functional InCA to re-introduce relations.

6.1 Computing a control-flow graph functionally

Consider we want to compute the control-flow graph (CFG) of a program as part of a Datalog-based program analysis. We want to represent the CFG such that it corresponds to a Datalog relation, so that we can easily compute its transitive closure later. While the functions of functional InCA compile to Datalog relations, our functions cannot be used to encode arbitrary relations. In particular, a function (`def flow(from: Stm): Stm = e`) cannot handle conditional statements that fork the control flow and connect to multiple successor statements. To support such relations, we must extend our frontend language.

We want to extend functional InCA in a way that integrates functions and relations elegantly. This is a language-design challenge and therefore naturally somewhat subjective. But it is the reason why we rejected the first idea that came to mind: to introduce relations next to functions. For example, a top-level relation (`rel flow(from: Stm, to: Stm) :- constraints`) could capture the CFG of a program. The problem is that we are now back at constraint programming, which is exactly what we wanted to avoid with functional InCA.

We propose a different extension of functional InCA that not only avoids this problem but that is simpler too: We introduce sets and tuples. Immutable sets and tuples are staple ingredients of functional programming and programmers already know how to use them. Moreover, any relation can be encoded as a set containing tuples of related values. Thus, the only question is if and how we can map functional programs over sets and tuples to Datalog. But first, let us illustrate how the extended functional InCA can be used.

In their classic textbook, Nielson et al. [16] compute the control flow of a `While`-statement through three functions. We can represent these functions in the extended functional InCA almost verbatim as shown in Figure 11. Here, `init` is a regular function whereas `final` and `flow` compute sets. A set literal `{e1, ..., en}` constructs a set and set union `++` composes two sets. For example, `final` uses these features to compute the final statement of each conditional branch. Sets can be processed through set comprehensions as shown in the definition of `flow` which can be seen in Figure 1. In particular, `(x1, ..., xn) in set` retrieves the elements of `set`, binds those `x` that are free, and tests for membership of those `x` that are bound.

Our encoding of relations makes it easy to implement computations that exercise Datalog's declarative fixpoint semantics, such as transitive closure, cycle detection, and recursive aggregation. We have already demonstrated such computations in the introduction of this paper and refrain from repeating them here. Instead, we show how to translate functional programs with sets and tuples to Datalog.

(functions) $F ::= \dots \mid [\text{@main}] \text{ def } f(\overline{x:T}) : \text{Set}[T] = s$
 (set expressions) $s ::= \{\overline{e}\} \mid s ++ s \mid \{e \mid \overline{pred}\} \mid \text{let } x = e \text{ in } s \mid \text{if } (e) s \text{ else } s \mid f(\overline{e})$
 (predicates) $pred ::= e \mid e \text{ in } s \mid e \text{ in } N$
 (expressions) $e ::= \dots \mid \text{fold}(f, f, f)$

■ **Figure 12** Extended abstract syntax with set and set operations.

$$\begin{aligned}
 \llbracket \{\overline{e}\} \rrbracket &= \bigcup_{e \in \overline{e}} \llbracket e \rrbracket \\
 \llbracket s_1 ++ s_2 \rrbracket &= \llbracket s_1 \rrbracket \cup \llbracket s_2 \rrbracket \\
 \llbracket \{e \mid p_1, \dots, p_n\} \rrbracket &= \{(t, \{t_1 = \text{true}, \dots, t_n = \text{true}\} \cup a \cup a_1 \cup \dots \cup a_n) \\
 &\quad \mid (t, a) \in \llbracket e \rrbracket, (t_1, a_1) \in \llbracket p_1 \rrbracket_{pred}, \dots, (t_n, a_n) \in \llbracket p_n \rrbracket_{pred}\} \\
 \llbracket \text{fold}(f_{init}, f_{op}, f_{set}) \rrbracket &= \{(\text{aggregate}(f_{set}, \text{toPrimitiveFun}(f_{init}), \text{toPrimitiveFun}(f_{op})), \emptyset)\} \\
 \llbracket e \rrbracket_{pred} &= \llbracket e \rrbracket \\
 \llbracket e \text{ in } s \rrbracket_{pred} &= \{(\text{true}, \{t_1 = t_2\} \cup a_1 \cup a_2 \mid (t_1, a_1) \in \llbracket e \rrbracket, (t_2, a_2) \in \llbracket s \rrbracket)\} \\
 \llbracket e \text{ in } N \rrbracket_{pred} &= \{(\text{true}, \{N(t)\} \cup a \mid (t, a) \in \llbracket e \rrbracket)\}
 \end{aligned}$$

■ **Figure 13** Compiling sets and set operations to Datalog.

6.2 Translating tuples and first-order sets to Datalog

The translation of sets and tuples to Datalog is mostly straightforward except for one thing: neither sets nor tuples are first-class in Datalog. For tuples this is hardly an issue since we can simply flatten tuples when translating them to Datalog. For example, a function $f_{oo}(\tau : (T_1, \dots, T_n)) : (U_1, \dots, U_m)$ becomes a flat relation $f_{oo}(T_1, \dots, T_n, U_1, \dots, U_m)$, and a function call $f_{oo}(e)$ becomes $f_{oo}(t_1, \dots, t_n, u_1, \dots, u_m)$, where e translates to n terms (t_1, \dots, t_n) and the function call yields m result terms (u_1, \dots, u_m) . Although our implementation supports tuples, we omit tuples from our translation semantics and focus on sets instead.

We want to translate sets to Datalog relations, but relations are first-order in Datalog and can only appear as top-level definitions. Thus, if we want to support first-class sets in functional IncA, we need to lift those sets first. For example, to translate a call `transitive` $\{(1,2), (2,3), (3,4)\}$ to Datalog, we have to translate $\{(1,2), (2,3), (3,4)\}$ to a top-level relation that can be queried from within `transitive`. To achieve this, we propose a clean solution in two steps:

1. We extend functional IncA first-order sets, which may only appear as function results. First-order sets translate to first-order relations as shown in the present subsection.
2. The subsequent subsection shows that a standard defunctionalization transformation simultaneously adds support for first-class functions and first-class sets to functional IncA.

Figure 12 defines the extended functional IncA, where we introduce first-order sets syntactically through a new non-terminal `s`. This syntactic differentiation does not replace type checking of the functional code, but serves to explain which expressions may yield sets without presenting functional IncA's type system, which is completely standard and uninteresting. First-order sets may only occur as the body of a function that yields a set and within other set expressions. A set comprehension can use predicates `pred` to check a boolean condition `e`, to query another set `(e in s)`, or to query all instances of an algebraic data type `(e in N)`. Here we finally see why we introduced instance relations for algebraic data types in Section 4. At last, we can convert a set to an atomic value through `fold(finit, fop, fset)`, where `fset` must be the name of a top-level definition.

$$\begin{array}{l}
\text{(expressions)} \quad e ::= \dots \mid f \mid (\overline{x:T}) \Rightarrow e \mid (\overline{x:T}) \Rightarrow s \mid e(\overline{e}) \\
\text{(types)} \quad T ::= \dots \mid \overline{T} \Rightarrow T
\end{array}$$

■ **Figure 14** Adding first-class functions to our Datalog frontend.

700 We extend $\llbracket \cdot \rrbracket$ to also handle set expressions s , and we add a translation function $\llbracket \cdot \rrbracket_{pred}$
701 for predicates. Figure 13 shows both translation functions. A set literal translates to a set of
702 alternative terms and set union computes the union of alternative terms. A set comprehension
703 builds all terms t generated by e for which all predicates are **true**.

We can only translate folds if the targeted Datalog engine supports aggregation over user-defined functions. In our experience, such user-defined functions must be implemented in the same language as the Datalog engine (e.g., C++ for Soufflé, a JVM language for Formulog and IncA). Thus, fold operations are considered built-in functions φ by Datalog engines. We extend the Datalog IR with aggregation accordingly:

$$\text{(Datalog terms)} \quad t ::= \dots \mid \text{aggregate}(R, \varphi, \varphi)$$

704 All we have left to do is to translate frontend functions f to built-in functions φ , which we
705 assume function `toPrimitiveFun` accomplishes. In our implementation, we target IncA and
706 compile user-defined frontend functions to Scala, which was straightforward. Soufflé does
707 not support aggregation over user-defined functions, hence we cannot target Soufflé if the
708 functional IncA program contains fold operations.

709 6.3 First-class functions and first-class sets

710 Functional IncA paves the road for transferring insights from functional programming
711 languages to Datalog. Here, we exemplify this potential by studying defunctionalization in
712 the context of functional IncA. Defunctionalization [18] is a well-known compilation technique
713 that compiles higher-order functions into first-order functions and first-class function values
714 into algebraic data. In particular, defunctionalization generates auxiliary *apply* functions that
715 dispatch on the algebraic data to execute the corresponding function body. Since functional
716 IncA supports first-order functions and algebraic data types, we can apply defunctionalization
717 to extend functional IncA with first-class functions.

718 Figure 14 shows how we extend functional IncA’s syntax with first-class functions. A
719 function value is either a reference to a top-level function f or a lambda. Note that we permit
720 lambdas to yield sets, since they will translate to first-order functions, which we translate
721 to first-order relations. Finally, we adapt function application to allow any expressions in
722 function position.

723 For example, consider an excerpt from our data-flow analyses of the While language:

```

724 def findExps(exp: Exp, f: Exp => Boolean): Set[Exp] = (exp match {
725   case Var(s) => {}
726   case Num(i) => {}
727   case Add(e1, e2) => findExps(e1, f) ++ findExps(e2, f)
728 }) ++ (if (f(exp)) {exp} else {})
729 def freevars(exp: Exp): Set[String] = {varName(e) | e in findExps(exp, isVar)}
730 def availableExps(exp: Exp): Set[Exp] = findExps(exp, (e: Exp) => e match {
731   case Var(s) => false
732   case Num(i) => false
733   case Add(e1, e2) => true })
734
735 
```

736 We define a higher-order function `findExps` that selects all subexpressions satisfying predicate
737 f . We use `findExps` twice, once to find all free variables of an expression and once to find all
738 non-trivial subexpressions. We implement a standard defunctionalization transformation
739 that translates this program into a first-order functional program:

```

740
741 data Defun0 = Funref0() | Lambda0()
742 def applyDefun0(fun: Defun0, e: Exp): Boolean = fun match {
743   case Funref0() => isVar(e)
744   case Lambda0() => e match {
745     case Var(s) => false
746     case Num(i) => false
747     case Add(e1, e2) => true } }
748 def findExps(exp: Exp, f: Defun0): Set[Exp] = (...) ++ (if (applyDefun0(f, exp)) {exp}
749   else {})
750 def freevars(exp: Exp): Set[String] = {varName(e) | e in findExps(exp, Funref0())}
751 def availableExps(exp: Exp): Set[Exp] = findExps(exp, Lambda0())
752

```

753 We can then translate the defunctionalized program to Datalog as described before. Thus,
 754 we have successfully extended functional IncA with first-class functions.

755 But how does this enable first-class sets? We already added support for first-order sets,
 756 which may only occur as function results. But since first-class functions translate to first-order
 757 functions, first-class functions may also yield sets. Thus, we can encode a first-class set *s* as
 758 a thunk `() => s`. For example, we can define a higher-order relation `transitive` as follows:

```

759
760 def transitive(cfg: () => Set[(Stm, Stm)]): Set[(Stm, Stm)] =
761   cfg() ++ {(s1,s3) | (s1,s2) in cfg(), (s2,s3) in transitive(cfg)}
762 @main def transitiveFlow(prog: Stm): Set[(Stm, Stm)] = let cfg = () => flow(prog) in
763   transitive(cfg)
764

```

765 Since function values become algebraic data, thunk-encoded sets are truly first-class: They
 766 can be assigned to variables and they can be passed as arguments. This shows how functional
 767 IncA permits insights from functional programming languages to carry over to Datalog,
 768 where they can unleash additional benefits.

769 **7 Case Studies: Data-Flow Analyses and Clone Detection**

770 We have presented functional Datalog frontend with relations that compiles to Datalog. In
 771 these final case studies, we want to demonstrate why this design is useful and how it enables
 772 a new way of implementing Datalog-based static analyses. To this end, we implemented
 773 flow-sensitive reaching definitions and interval analyses for the WHILE language in functional
 774 IncA. Additionally, we show how to describe clone detection of Java bytecode.

775 **7.1 Data-Flow Analyses**

776 Figure 15 shows an excerpt of the reaching definitions analysis, which determines where a
 777 variable was last defined. Our analysis implementation is completely standard except that
 778 we use a `retain` filter in place of the usual `kill` set. This is because functional IncA does
 779 not support negation yet, which is needed for set difference. We hope to extend functional
 780 IncA with negation in future work, but note that negation in Datalog is far from trivial and
 781 deserves a separate study.

782 The reaching definitions case study shows how we benefit from using functions and
 783 relations. The main benefit of functions is the ease of implementation in a well-known
 784 programming paradigm, as illustrated by `gen` in our example. The main benefit of relations is
 785 the implicit fixpoint semantics provided by Datalog. Specifically, note that `entry` and `exit` call
 786 each other unconditionally and diverges under functional-programming semantics. However,
 787 Datalog implicitly computes the least fixpoint of relations, which is computable because the
 788 relations are finite: There are only finitely many variables and assignments in any program.
 789 Here, functional IncA reaps the rewards of compiling to Datalog.


```

def gen(stm: Stm): Set[(String,Maybe[Stm])] = stm match {
  case Assign(x, a) => {(x, Just(stm))}
  case Sequence(s1, s2) => {}
  case If(c, s1, s2) => {}
  case While(c, s) => {} }
def retain(stm: Stm, x: String): Boolean = ...
def entry(stm: Stm, prog: Stm): Set[(String, Maybe[Stm])] =
  if (stm == init(prog)) {(x, Nothing()) | x in freevarsStm(prog)}
  else {(x,d) | (pred, stm) in flow(prog), (x,d) in exit(pred, prog)}
def exit(stm: Stm, prog: Stm): Set[(String,Maybe[Stm])] =
  gen(stm) ++ {(x,d) | (x,d) in entry(stm, prog), retain(stm, x)}
@main def allExits(prog: Stm): Set[(Stm, String, Maybe[Stm])] =
  {(s,x,d) | s in Stm, (x,d) in exit(s, prog)}

```

■ **Figure 15** A reaching definitions analysis for the While language that we compile to Datalog.

```

data Val = BotVal() | IntervalVal(Interval) | BoolVal(Bool) | TopVal()
...
// entry_var, exit_var, and aeval as shown in Figure 2 (Introduction)
def add(v1: Val, v2: Val): Val = ...
def addInterval(iv1: Interval, iv2: Interval): Interval = iv1 match {
  case TopInterval() => TopInterval()
  case IV(l1, h1) => iv2 match {
    case TopInterval() => TopInterval()
    case IV(l2, h2) => IV(l1 + l2, h1 + h2) } }
def joinVal(v1: Val, v2: Val): Val = ...
def joinInterval(iv1: Interval, iv2: Interval): Interval = iv1 match {
  case TopInterval() => TopInterval()
  case IV(l1, h1) => iv2 match {
    case TopInterval() => TopInterval()
    case IV(l2, h2) => widenInterval(IV(Math.min(l1, l2), Math.max(h1, h2))) } }

```

■ **Figure 16** Interval analysis of the While language using abstract interpretation.

790 In the reaching definitions analysis, the fixpoint computation within `entry` and `exit` only
791 invokes simple functions `gen` and `retain`. Therefore, it is reasonable to implement the reaching
792 definitions in Datalog directly, although we believe functional IncA is easier to use. In
793 contrast, our second data-flow analysis implements an interval analysis that requires complex
794 functions to abstractly interpret expressions. We show an excerpt of the interval analysis in
795 Figure 16 and Figure 2. We use data type `val` to represent abstract values and use relations
796 `entry_var` and `exit_var` to map variables to their abstract value. For an `Assign` statement,
797 `exit_var` invokes an abstract interpreter `aeval` that computes the abstract value of the assigned
798 expression. Even for this simple WHILE language, the abstract interpreter already consists
799 of 90 lines of functional code that compile to 342 lines of Datalog code. Moreover, `aeval` is
800 part of the fixpoint loop, because it invokes `entry_var` for variable references, which invokes
801 `exit_var`, which invokes `aeval`. Therefore, `aeval` really must translate to Datalog rules and
802 cannot be represented as a built-in function, because then it could not invoke `entry_var`.
803 Finally, note that we use a user-defined function `joinVal` to aggregate abstract values in
804 `entry_var`. In particular, `joinVal` implements widening on intervals to ensure the analysis
805 always terminates. All of these concerns are easy to address in functional IncA, because we
806 can use functional programming while relying on Datalog’s fixpoint semantics.

807 7.2 Clone Detection

808 Figure 17 shows an excerpt of how to construct an abstract syntax tree of Shimple code
809 and apply clone-detection techniques such as testing for alpha-equivalence. Shimple is

```

def getStm(inst: Instruction): Set[Stm] = {
  InvokeStm(recvExp, meth, args) |
  (inst, v) not in _AssignReturnValue,
  (inst, _, meth, recv, _) in _VirtualMethodInvocation,
  recvExp in getExp(recv),
  args in getArgs(inst, 0) } ++ ...
def getExp(v: String): Set[Exp] = ...
def getArgs(inst: Instruction, currentIdx: Int): Set[List[Exp]] = ...
def isStmClone(s1:Stm, s2:Stm, iPairs:NPairs, lPairs:NPairs): Boolean = (s1,s2) match {
  case (InvokeStm(r1, m1, a1), InvokeStm(r2, m2, a2)) =>
    m1 == m2 && isExpClone(r1, r2, iPairs) && isArgListClone(a1, a2, iPairs)
  ... }

```

■ **Figure 17** Clone detection of Shimple Code.

810 a variant of the Java bytecode representation Jimple [28] in SSA form. To access the
 811 Shimple representation, we extend functional IncA to read Soufflé relations, because the
 812 Doop framework [10] generates Soufflé facts. Using Soufflé facts enables us to detect clones
 813 of real-world Java programs. The Soufflé relations are prefixed by an underscore. Technically,
 814 we compile the Soufflé program and the functional program to a single Datalog program.
 815 However, we do not derive demand patterns for relations of the Soufflé program.

816 The function `getStm` constructs an abstract syntax tree representation of Shimple code. We
 817 highlight the case for constructing an invocation statement. We only generate an invocation
 818 statement for an instruction `inst` if it is a virtual method invocation and the instruction does
 819 not assign a return value for the given instruction. Note that functional IncA does not allow
 820 negation in general. However, it is possible to query Soufflé relations negatively as we do not
 821 apply the demand transformation to Soufflé relations. Hence, the demand transformation
 822 does not introduce negated dependency cycles. We generate the receiver of the method
 823 call by using function `getExp` which constructs an expression tree given a variable name. At
 824 last, we construct the argument of the given invocation statement by calling `getArgs`. Note
 825 while `getStm`, `getExp` and `getArgs` have `Set` as return type, the functions yield singleton sets.
 826 Returning a set is necessary due to the fact that we query Soufflé relations.

827 Next, we use the constructed abstract syntax trees as a basis to detect clones. We show a
 828 clone-detection function `isStmClone` which checks if the statements are alpha-equivalent. We
 829 traverse the statements `s1` and `s2` simultaneously while checking that the statements and
 830 inner expressions are equal. Because we rely on Soufflé relations generated by the Doop
 831 framework [10], we could integrate static analysis information such as points-to information
 832 into clone detection. The case study shows that describing alpha-equivalence of Java bytecode
 833 in a functional style is straightforward. It is possible to realize more sophisticated clone-
 834 detection techniques using functional IncA such as structural diffing [11].

835 **8 Implementation and Performance Evaluation**

836 In this section we will discuss our implementation and do an early performance evaluation.

837 **8.1 Implementation**

838 We implemented functional IncA by compiling it to a Datalog IR provided by the IncA
 839 framework. The Datalog IR can target two different backends namely IncA and Soufflé without
 840 any change to the underlying Datalog solvers VIATRA [29] and Soufflé [20] respectively. Our
 841 compiler generates Datalog code as shown in this paper, including the demand transformation.
 842 This implementation not only demonstrates the feasibility of our design, but also shows how

843 advantageous it is to reuse existing Datalog solvers. In particular, the VIATRA Datalog solver
 844 supports incrementality: Changes in extensional relations trigger incremental updates in
 845 derived relations. We inherit this incrementality for free. For example, we can run the interval
 846 analysis of Subsection 7.1 incrementally by diffing the input programs and feeding the resulting
 847 patch to IncA [11]. Targeting Soufflé allows us to generate efficient and scalable C++ programs
 848 that run on multi-core machines. However, Soufflé does not support user-defined aggregation,
 849 hence we do not support translating functional IncA programs containing fold operations.
 850 The implementation is available at <https://gitlab.rlp.net/plmz/inca-scala>.

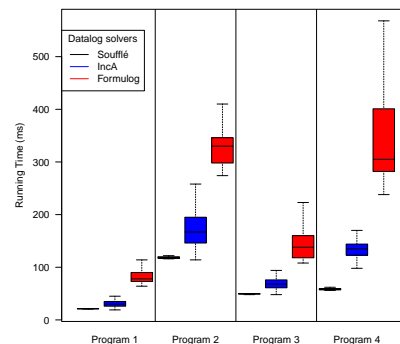
851 8.2 Performance Evaluation

852 We evaluate the performance of functional IncA and show that it is advantageous to use
 853 established solvers instead of implementing custom Datalog solvers for new frontends. We
 854 compare the running times of executing a data-flow analysis for the While language run with
 855 Soufflé, IncA, and Formulog. We choose Soufflé and IncA as they are already established
 856 Datalog frameworks. We choose Formulog because it is one representative of the *frontend-first*
 857 approach which combines first-order ML functions with Datalog by implementing a custom
 858 Datalog solver. Even though IncA uses an incremental Datalog solver VIATRA [29], we do
 859 not measure the incremental performance of IncA which we leave as future work.

860 The data-flow analysis that we run is an adapted interval analysis. The analysis collects
 861 all integers $-100 \leq i \leq 100$, a variable can be assigned to. Whenever we encounter an integer
 862 $i < -100$ we return the default value -1000 and when we encounter an integer $i > 100$ we
 863 return 1000 . We implement this cut-off to ensure that the data-flow analysis terminates in
 864 the presence of loops. We have chosen this type of analysis instead of an interval analysis,
 865 because an interval analysis requires user-defined aggregation which Formulog and Soufflé
 866 currently do not support. We implement four different programs as input of the data-flow
 867 analysis. The programs consist of nested *while* and *if* statements and are designed in such a
 868 way that a lot of information has to be propagated along the edges of the control-flow graph.

869 For Formulog and IncA, both of which are Datalog solvers that run on the JVM, we
 870 first do 10 warmup runs and then measure 90 runs. We do not measure the time it takes
 871 to initialize the extensional database but only measure the running times of deriving the
 872 intensional database. For Soufflé, we compile an executable and measure the running time of
 873 the compiled Soufflé solver to derive the intensional database 9 times. Note that we do not
 874 warmup for Soufflé programs as they are compiled to C++ and then to executable machine
 875 code. We store the extensional database within input files and do not measure the I/O
 876 actions needed to read those input files. We load the contents of the input files into RAM by
 877 executing the compiled Soufflé program once. Hence, the following measured runs access the
 878 extensional database stored in RAM. We performed our benchmarks on a machine with an
 879 Intel Core i7 at 2.7 GHz with 16 GB of RAM, running 64-bit OSX 11.4, Java 1.14.0_1.

880 We show the running times of deriving the intensional database in milliseconds for each program on the
 881 right-hand side. We see that the custom Datalog solver for Formulog is slower than the established solvers
 882 such as Soufflé and IncA for all input programs. The Formulog solver is $\sim 3.7x$ slower than the Soufflé solver
 883 and $\sim 2.2x$ slower than the IncA solver. Note that the compiled executable of Soufflé has the fastest running
 884 time of all three solvers. This shows that it is desirable to compile Datalog with functional constructs to already
 885
 886
 887
 888
 889



890 established Datalog dialects instead of implementing
891 custom solvers for new Datalog frontends if possible.

892 **9 Related Work**

893 We propose functional programming with sets as a frontend for Datalog to replace Datalog's
894 traditional constraint programming. This design differs from most prior works, which retain
895 constraint programming as a basis and add functional aspects on top of it. Our approach has
896 three advantages: (i) functional programming is easy to use, (ii) we can compute fixpoints
897 across functions and relations, and (iii) we can reuse existing Datalog solvers. In the remainder
898 of this section, we discuss related work.

899 IncA is an incremental Datalog framework that supports recursive aggregation over
900 user-defined functions and data types [22, 23]. These user-defined functions and data types
901 must be implemented in a JVM language and cannot query Datalog relations. The original
902 frontend of IncA provides a shallow abstraction over Datalog called *pattern functions* [24].
903 These pattern functions consist of sequences of constraints and really are not comparable to
904 the functional programming we support in functional IncA.

905 Flix [14] exposes constraint programming to the user, but extends it with functional
906 programming. The runtime system of Flix executes functional code but also contains a custom
907 Datalog solver. While functional and Datalog aspects are intertwined in Flix, they cannot
908 interact as tightly as the functions and relations in our approach. Specifically, user-defined
909 functions cannot recursively query derived relations, as required by our interval analysis.
910 However, it is also not obvious how to extend our approach to compile Flix to Datalog,
911 because Flix supports the generation of additional Datalog constraints at run time [13].

912 Formulog [9] combines first-order ML functions with Datalog and SMT solvers. In
913 particular, Datalog rules can contain ML expressions and ML code can recursively query
914 Datalog relations. Formulog's runtime understands both languages, which is why a custom
915 Datalog solver was needed that can evaluate ML expressions and Datalog constraints
916 interleaved. Our approach should naturally extend to Formulog. Indeed, we could add SMT
917 solving as a built-in function (`def solveSMT(spec: String): String`) and rely on user-defined
918 data types for SMT formulae and models, both of which are built-in types in Formulog.

919 Datafun [6] defines a higher-order functional programming language with sets and fixpoint
920 semantics. From a language-design perspective, Datafun is the most closely related work.
921 Both languages support commonly known functional expressions. One difference is how
922 fixpoint computations are expressed in the surface syntax. Datafun provides a fixpoint
923 expression which explicitly states over which function a fixpoint will be computed. However,
924 in functional IncA the fixpoint computation is not explicitly given but implicitly given by
925 the dependencies between functions. Datafun and functional IncA follow different design
926 philosophies. Datafun provides a termination guarantee: If a Datafun program is well-typed,
927 then a unique least fixed point exists and the program will terminate. Our language does
928 not provide such a guarantee since a well-typed program can still diverge. Consequently,
929 Datafun is more restrictive to guarantee termination while functional IncA gives developers
930 more freedom (and responsibility). Datafun requires that the lattice type over which a
931 fixpoint is computed does not contain an infinite ascending chain. One disadvantage of
932 Datafun's design is that some programs that terminate in our system are not accepted by
933 Datafun. For example, the interval analysis we presented is not well-typed in Datafun as the
934 interval lattice has an infinite ascending chains. To ensure termination in functional IncA,
935 we had to introduce widening to break the infinite ascending chains of the interval lattice.
936 Many interesting static analyses use infinite lattices with infinite ascending chains. Hence,

937 Datafun cannot be used to express such analyses. While it is an interesting question how
938 to guarantee termination for as many programs as possible, our system is more viable to
939 implement real-world programs. Another difference is that Datafun has its own bottom-up
940 semantics which was recently extended to support semi-naïve evaluation [5]. In contrast,
941 we translate programs to Datalog and utilize off-the-shelf Datalog solvers, which readily
942 implement semi-naïve evaluation and other optimizations.

943 QL [7] is a logic programming language with object-oriented features such as classes and
944 methods to structure logic programs. QL compiles to Datalog to encode inheritance and
945 virtual dispatch of member predicates. We also propose to compile to Datalog, but focus on
946 functional programming with algebraic data types. It would be interesting to see how we
947 can extend functional IncA with object-oriented features and how these interact.

948 Mercury [21] is a logic programming language that consists of relations and rules deriving
949 those relations. Like any Datalog, Mercury also supports the encoding of functions as relations,
950 but in Mercury users can additionally annotate parameters as inputs and deterministic
951 outputs. Mercury implements a custom Datalog solver that exploits such functional relations
952 by executing them like a deterministic program. It would be interesting to explore *generic*
953 Datalog optimizations that exploit functional relations, since we can easily generate the
954 necessary annotations in functional IncA.

955 Bloom [2, 3] is a domain-specific language for distributed systems that uses the Datalog
956 variant Dedalus [4] under the hood. Bloom provides built-in higher-order functions such as
957 `map` and `reduce` that operate over collections. Bloom is embedded in Ruby and user-defined
958 functions and data types can be written in Ruby, but these user-defined functions cannot
959 access the contents of relations. Therefore, we cannot describe an interval analysis in the
960 same style we have shown in the previous section in Bloom.

961 Soufflé [20] is an efficient Datalog solver that can interpret Datalog rules directly or translate
962 them to C++. It is possible to define user-defined functions as C++ functions, but again
963 these functions cannot access the contents of relations. Soufflé has support for algebraic data
964 types, but developers have to ensure that only finitely many values are constructed. For our
965 use cases, this amounts to encoding the input relations by hand.

966 **10 Conclusion**

967 Datalog is supposedly declarative, but many programs are hard to express as constraints.
968 We propose functional programming with sets as a new frontend for Datalog that solves this
969 problem: *functional IncA*. Specifically, we translate functional IncA programs to Datalog
970 and employ a demand transformation to ensure the Datalog program terminates whenever
971 the original program terminates. While users of functional IncA only need to learn a single
972 functional programming language, they enjoy Datalog's fixpoint semantics across functions
973 and relations. Moreover, since all generated code is pure Datalog, we can use off-the-shelf
974 Datalog solvers rather than building our own. Specifically, we implemented our approach as
975 a frontend for IncA [23] as well as Soufflé [20] and demonstrated how easy it is to express
976 complex Datalog programs with it. Our case studies include clone detection of real-world
977 Java programs, program analyses, a program transformation, and an interpreter, all of which
978 are easy to express functionally but translate to highly complex Datalog code. We have
979 shown through early performance measurements that it is indeed desirable to use established
980 Datalog solvers than implement custom solvers that embed a functional programming
981 language as Formlog did. In future work, we want to investigate the performance of the
982 generated Datalog code and study how compiler optimization can help. We also want to
983 support negation in functional IncA, but the demand transformation potentially breaks the

984 stratifiability of programs. We want to explore if the solution by Tekle and Liu [26] can be
 985 used. At last, we want to investigate how to properly debug functional InCA programs.

986 ——— References ———

- 987 **1** Serge Abiteboul, Zoë Abrams, Stefan Haar, and Tova Milo. Diagnosis of asynchronous
 988 discrete event systems: datalog to the rescue! In Chen Li, editor, *Proceedings of the Twenty-*
 989 *fourth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems,*
 990 *June 13-15, 2005, Baltimore, Maryland, USA*, pages 358–367. ACM, 2005. URL: <https://doi.org/10.1145/1065167.1065214>, doi:10.1145/1065167.1065214.
- 992 **2** Peter Alvaro, Tyson Condie, Neil Conway, Khaled Elmeleegy, Joseph M. Hellerstein, and
 993 Russell Sears. Boom analytics: exploring data-centric, declarative programming for the cloud.
 994 In Christine Morin and Gilles Muller, editors, *European Conference on Computer Systems,*
 995 *Proceedings of the 5th European conference on Computer systems, EuroSys 2010, Paris, France,*
 996 *April 13-16, 2010*, pages 223–236. ACM, 2010. URL: <https://doi.org/10.1145/1755913.1755937>, doi:10.1145/1755913.1755937.
- 998 **3** Peter Alvaro, Neil Conway, Joseph M. Hellerstein, and William R. Marczak. Consistency
 999 analysis in bloom: a CALM and collected approach. In *CIDR 2011, Fifth Biennial Conference*
 1000 *on Innovative Data Systems Research, Asilomar, CA, USA, January 9-12, 2011, Online*
 1001 *Proceedings*, pages 249–260. www.cidrdb.org, 2011. URL: [http://cidrdb.org/cidr2011/](http://cidrdb.org/cidr2011/Papers/CIDR11_Paper35.pdf)
 1002 [Papers/CIDR11_Paper35.pdf](http://cidrdb.org/cidr2011/Papers/CIDR11_Paper35.pdf).
- 1003 **4** Peter Alvaro, William R. Marczak, Neil Conway, Joseph M. Hellerstein, David Maier, and
 1004 Russell Sears. Dedalus: Datalog in time and space. In Oege de Moor, Georg Gottlob, Tim
 1005 Furche, and Andrew Jon Sellers, editors, *Datalog Reloaded - First International Workshop,*
 1006 *Datalog 2010, Oxford, UK, March 16-19, 2010. Revised Selected Papers*, volume 6702 of *Lecture*
 1007 *Notes in Computer Science*, pages 262–281. Springer, 2010. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-642-24206-9_16)
 1008 [978-3-642-24206-9_16](https://doi.org/10.1007/978-3-642-24206-9_16), doi:10.1007/978-3-642-24206-9_16.
- 1009 **5** Michael Arntzenius and Neel Krishnaswami. Seminaïve evaluation for a higher-order functional
 1010 language. *Proc. ACM Program. Lang.*, 4(POPL):22:1–22:28, 2020. URL: [https://doi.org/](https://doi.org/10.1145/3371090)
 1011 [10.1145/3371090](https://doi.org/10.1145/3371090), doi:10.1145/3371090.
- 1012 **6** Michael Arntzenius and Neelakantan R. Krishnaswami. Datafun: A functional Datalog. In
 1013 Jacques Garrigue, Gabriele Keller, and Eijiro Sumii, editors, *Proceedings of the 21st ACM*
 1014 *SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan,*
 1015 *September 18-22, 2016*, pages 214–227. ACM, 2016. URL: [https://doi.org/10.1145/2951913.](https://doi.org/10.1145/2951913.2951948)
 1016 [2951948](https://doi.org/10.1145/2951913.2951948), doi:10.1145/2951913.2951948.
- 1017 **7** Pavel Avgustinov, Oege de Moor, Michael Peyton Jones, and Max Schäfer. QL: object-
 1018 oriented queries on relational data. In Shriram Krishnamurthi and Benjamin S. Lerner,
 1019 editors, *30th European Conference on Object-Oriented Programming, ECOOP 2016, July*
 1020 *18-22, 2016, Rome, Italy*, volume 56 of *LIPICs*, pages 2:1–2:25. Schloss Dagstuhl - Leibniz-
 1021 Zentrum für Informatik, 2016. URL: <https://doi.org/10.4230/LIPICs.ECOOP.2016.2>, doi:
 1022 [10.4230/LIPICs.ECOOP.2016.2](https://doi.org/10.4230/LIPICs.ECOOP.2016.2).
- 1023 **8** Catriel Beeri and Raghu Ramakrishnan. On the power of magic. *The Journal of Logic*
 1024 *Programming*, 10(3):255–299, 1991. Special Issue: Database Logic Programming. URL:
 1025 <https://www.sciencedirect.com/science/article/pii/074310669190038Q>, doi:[https://](https://doi.org/10.1016/0743-1066(91)90038-Q)
 1026 [doi.org/10.1016/0743-1066\(91\)90038-Q](https://doi.org/10.1016/0743-1066(91)90038-Q).
- 1027 **9** Aaron Bembenek, Michael Greenberg, and Stephen Chong. Formulog: Datalog for SMT-
 1028 based static analysis. *Proc. ACM Program. Lang.*, 4(OOPSLA):141:1–141:31, 2020. URL:
 1029 <https://doi.org/10.1145/3428209>, doi:10.1145/3428209.
- 1030 **10** Martin Bravenboer and Yannis Smaragdakis. Strictly declarative specification of sophisticated
 1031 points-to analyses. In Shail Arora and Gary T. Leavens, editors, *Proceedings of the 24th*
 1032 *Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages,*

- 1033 and Applications, *OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*, pages 243–
1034 262. ACM, 2009. URL: <https://doi.org/10.1145/1640089.1640108>, doi:10.1145/1640089.
1035 1640108.
- 1036 11 Sebastian Erdweg, Tamás Szabó, and André Pacak and. Concise, type-safe, and efficient
1037 structural diffing. In *Programming Language Design and Implementation (PLDI)*. ACM, 2021.
- 1038 12 Shan Shan Huang, Todd Jeffrey Green, and Boon Thau Loo. Datalog and emerging applications:
1039 an interactive tutorial. In Timos K. Sellis, Renée J. Miller, Anastasios Kementsietsidis, and
1040 Yannis Velegarakis, editors, *Proceedings of the ACM SIGMOD International Conference on Man-*
1041 *agement of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 1213–1216. ACM,
1042 2011. URL: <https://doi.org/10.1145/1989323.1989456>, doi:10.1145/1989323.1989456.
- 1043 13 Magnus Madsen and Ondrej Lhoták. Fixpoints for the masses: programming with first-class
1044 Datalog constraints. *Proc. ACM Program. Lang.*, 4(OOPSLA):125:1–125:28, 2020. URL:
1045 <https://doi.org/10.1145/3428193>, doi:10.1145/3428193.
- 1046 14 Magnus Madsen, Ming-Ho Yee, and Ondrej Lhoták. From Datalog to Flix: A declarative
1047 language for fixed points on lattices. In Chandra Krintz and Emery Berger, editors, *Proceedings*
1048 *of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation,*
1049 *PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, pages 194–208. ACM, 2016. URL:
1050 <https://doi.org/10.1145/2908080.2908096>, doi:10.1145/2908080.2908096.
- 1051 15 David Maier, K. Tuncay Tekle, Michael Kifer, and David Scott Warren. Datalog: concepts,
1052 history, and outlook. In Michael Kifer and Yanhong Annie Liu, editors, *Declarative Logic*
1053 *Programming: Theory, Systems, and Applications*, pages 3–100. ACM / Morgan & Claypool,
1054 2018. URL: <https://doi.org/10.1145/3191315.3191317>, doi:10.1145/3191315.3191317.
- 1055 16 Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. *Principles of program analysis*.
1056 Springer, 1999.
- 1057 17 André Pacak, Sebastian Erdweg, and Tamás Szabó. A systematic approach to deriving
1058 incremental type checkers. *Proc. ACM Program. Lang.*, 4(OOPSLA):127:1–127:28, 2020. URL:
1059 <https://doi.org/10.1145/3428195>, doi:10.1145/3428195.
- 1060 18 John C. Reynolds. Definitional interpreters for higher-order programming languages. *High. Or-*
1061 *der Symb. Comput.*, 11(4):363–397, 1998. URL: <https://doi.org/10.1023/A:1010027404223>,
1062 doi:10.1023/A:1010027404223.
- 1063 19 Bernhard Scholz, Herbert Jordan, Pavle Subotic, and Till Westmann. On fast large-scale
1064 program analysis in Datalog. In Ayal Zaks and Manuel V. Hermenegildo, editors, *Proceedings*
1065 *of the 25th International Conference on Compiler Construction, CC 2016, Barcelona, Spain,*
1066 *March 12-18, 2016*, pages 196–206. ACM, 2016. URL: <https://doi.org/10.1145/2892208.2892226>,
1067 doi:10.1145/2892208.2892226.
- 1068 20 Bernhard Scholz, Kostyantyn Vorobyov, Padmanabhan Krishnan, and Till Westmann. A
1069 Datalog source-to-source translator for static program analysis: An experience report. In
1070 *24th Australasian Software Engineering Conference, ASWEC 2015, Adelaide, SA, Australia,*
1071 *September 28 - October 1, 2015*, pages 28–37. IEEE Computer Society, 2015. URL: <https://doi.org/10.1109/ASWEC.2015.15>,
1072 doi:10.1109/ASWEC.2015.15.
- 1073 21 Zoltan Somogyi, Fergus Henderson, and Thomas C. Conway. The execution algorithm of
1074 mercury, an efficient purely declarative logic programming language. *J. Log. Program.*, 29(1-
1075 3):17–64, 1996. URL: [https://doi.org/10.1016/S0743-1066\(96\)00068-4](https://doi.org/10.1016/S0743-1066(96)00068-4),
1076 doi:10.1016/S0743-1066(96)00068-4.
- 1077 22 Tamás Szabó, Gábor Bergmann, Sebastian Erdweg, and Markus Voelter. Incrementalizing
1078 lattice-based program analyses in Datalog. *Proc. ACM Program. Lang.*, 2(OOPSLA):139:1–
1079 139:29, 2018. URL: <https://doi.org/10.1145/3276509>, doi:10.1145/3276509.
- 1080 23 Tamás Szabó, Sebastian Erdweg, and Gábor Bergmann. Incremental whole-program analysis
1081 in Datalog with lattices. In *Programming Language Design and Implementation (PLDI)*. ACM,
1082 2021.
- 1083 24 Tamás Szabó, Sebastian Erdweg, and Markus Voelter. Inca: a DSL for the definition of
1084 incremental program analyses. In David Lo, Sven Apel, and Sarfraz Khurshid, editors,

- 1085 *Proceedings of the 31st IEEE/ACM International Conference on Automated Software En-*
1086 *gineering, ASE 2016, Singapore, September 3-7, 2016*, pages 320–331. ACM, 2016. URL: <https://doi.org/10.1145/2970276.2970298>, doi:10.1145/2970276.2970298.
- 1087
- 1088 **25** K. Tuncay Tekle and Yanhong A. Liu. Precise complexity analysis for efficient datalog queries.
1089 In Temur Kutsia, Wolfgang Schreiner, and Maribel Fernández, editors, *Proceedings of the*
1090 *12th International ACM SIGPLAN Conference on Principles and Practice of Declarative*
1091 *Programming, July 26-28, 2010, Hagenberg, Austria*, pages 35–44. ACM, 2010. URL: <https://doi.org/10.1145/1836089.1836094>, doi:10.1145/1836089.1836094.
- 1092
- 1093 **26** K. Tuncay Tekle and Yanhong A. Liu. Extended magic for negation: Efficient demand-driven
1094 evaluation of stratified Datalog with precise complexity guarantees. In Bart Bogaerts, Esra
1095 Erdem, Paul Fodor, Andrea Formisano, Giovambattista Ianni, Daniela Inclezan, Germán Vidal,
1096 Alicia Villanueva, Marina De Vos, and Fangkai Yang, editors, *Proceedings 35th International*
1097 *Conference on Logic Programming (Technical Communications), ICLP 2019 Technical Commu-*
1098 *nications, Las Cruces, NM, USA, September 20-25, 2019*, volume 306 of *EPTCS*, pages 241–254,
1099 2019. URL: <https://doi.org/10.4204/EPTCS.306.28>, doi:10.4204/EPTCS.306.28.
- 1100 **27** Jeffrey D. Ullman. Bottom-up beats top-down for Datalog. In Avi Silberschatz, editor,
1101 *Proceedings of the Eighth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of*
1102 *Database Systems, March 29-31, 1989, Philadelphia, Pennsylvania, USA*, pages 140–149. ACM
1103 Press, 1989. URL: <https://doi.org/10.1145/73721.73736>, doi:10.1145/73721.73736.
- 1104 **28** Raja Vallee-Rai and Laurie J Hendren. Jimple: Simplifying java bytecode for analyses and
1105 transformations. 1998.
- 1106 **29** Dániel Varró, Gábor Bergmann, Ábel Hegedüs, Ákos Horváth, István Ráth, and Zoltán
1107 Ujhelyi. Road to a reactive and incremental model transformation platform: three generations
1108 of the VIATRA framework. *Software & Systems Modeling*, 15(3):609–629, Jul 2016. URL:
1109 <https://doi.org/10.1007/s10270-016-0530-4>, doi:10.1007/s10270-016-0530-4.