

Integrating Logic Rules with Everything Else, Seamlessly

YANHONG A. LIU, SCOTT D. STOLLER, YI TONG and BO LIN

Stony Brook University, Stony Brook, NY 11794, USA

(e-mails: liu@cs.stonybrook.edu, stoller@cs.stonybrook.edu, yittong@cs.stonybrook.edu,
bolin@cs.stonybrook.edu)

submitted 6 February 2023; revised 7 April 2023; accepted 12 May 2023

Abstract

This paper presents a language, Alda, that supports all of logic rules, sets, functions, updates, and objects as seamlessly integrated built-ins. The key idea is to support predicates in rules as set-valued variables that can be used and updated in any scope, and support queries using rules as either explicit or implicit automatic calls to an inference function. We have defined a formal semantics of the language, implemented a prototype compiler that builds on an object-oriented language that supports concurrent and distributed programming and on an efficient logic rule system, and successfully used the language and implementation on benchmarks and problems from a wide variety of application domains. We describe the compilation method and results of experimental evaluation.

KEYWORDS: language design and implementation, logic rules, sets, comprehension, aggregation, quantification, functions, updates, objects, concurrent and distributed

1 Introduction

Logic rules are powerful for expressing complex reasoning and analysis problems, especially in critical areas such as program analysis, decision support, networking, and security (Warren and Liu 2017; Liu 2018). However, developing application programs that use logic rules remains challenging:

- Powerful logic languages and systems support succinct use of logic rules for complex reasoning and analysis, but not as directly or conveniently for many other aspects of applications—for example, data aggregation, numerical computation, input/output, modular construction, and concurrency—that are more easily expressed using set queries, functions, state updates, and object encapsulation (Maier *et al.* 2018).
- At the same time, commonly used languages for building applications support many powerful features but not logic rules, and to use a logic rule system, tedious and error-prone interface code is required—to pass rules and data to the rule system, invoke operations of the rule system for answering queries, and pass the results back—manually solving an impedance mismatch, similarly as in interfaces with relational databases (Geiger 1995), making logic rules harder to use than necessary.

What is lacking is (1) a simple and powerful language that can express application problems by directly using logic rules as well as all other features without extra interface

code, and with a clear semantics for analysis as well as execution, plus (2) a compilation framework for implementing this powerful language, in a practical way by extending a widely used programming language, and leveraging best performance of logic programming systems.

We have developed such a powerful language, Alda, that combines the advantages of logic languages and commonly used languages for building applications, by supporting direct use of all of logic rules, sets, functions, updates, and objects including concurrent and distributed processes as seamlessly integrated built-ins with no extra interfaces.

- Sets of rules can be specified directly as other definitions can, where predicates in rules are simply set-valued variables holding the set of tuples for which the predicate is true. Thus, predicates can be used directly as set-valued variables and vice versa without needing any extra interface, and predicates being set-valued variables are completely different from functions or procedures, unlike in prior logic rule languages and extensions.
- Queries using rule sets are calls to an inference function that computes desired values of derived predicates (i.e., predicates in conclusions of rules) given values of base predicates (i.e., predicates not in conclusions of rules). Thus, queries as function calls need no extra interface either, and a rule set can be used with predicates in it holding the values of any appropriate set-valued variables.
- Values of predicates can be updated either directly as for other variables or by the inference function; declarative semantics of rules are ensured by automatically maintaining values of derived predicates when values of base predicates are updated, through appropriate implicit calls to the inference function.
- Predicates and rule sets can be object attributes as well as global and local names, just as variables and functions can.

We also defined a formal semantics that integrates declarative and operational semantics. The integrated semantics supports, seamlessly, all of logic programming with rules, database programming with sets, functional programming, imperative programming, and object-oriented programming including concurrent and distributed programming. Note that predicates as variables, and queries as calls with different predicate values, also avoid the need for higher-order predicates or more sophisticated features for reusing rules on different predicates in more complex logic languages.

Implementing such a powerful language is nontrivial, especially to support logic rules together with updates and objects. We describe a compilation framework for implementation that achieves generally good performance.

- The framework implements Alda by building on an object-oriented language that supports all other features but not logic rules, and uses an efficient logic rule system for queries using rules.
- The framework considers and analyzes different kinds of updates to predicates in different scopes and uses an efficient implementation for each kind to minimize calls to the inference function while still ensuring the declarative semantics of rules.
- The framework also allows optimizations from decades of study of logic rules to be added for further efficiency improvements, both for queries using rules and for incremental queries under updates.

```

1 class CoreRBAC: # class for Core RBAC component/object
2   def setup(): # method to set up the object, with no arguments
3     self.USERS, self.ROLES, self.UR := {}, {}, {}
4     # set users, roles, user-role pairs to empty sets
5   def AddRole(role): # method to add a role
6     ROLES.add(role) # add the role to ROLES
7   def AssignedUsers(role): # method to return assigned users of a role
8     return {u: u in USERS | (u,role) in UR} # return set of users having the role
9   ...
9 class HierRBAC extends CoreRBAC: # Hierarchical RBAC extending Core RBAC
10  def setup():
11    super().setup() # call setup of CoreRBAC, to set sets as in there
12    self.RH := {} # set ascendant-descendant role pairs to empty set
13  def AddInheritance(a,d): # to add inherit. of an ascendant by a descendant
14    RH.add((a,d)) # add pair (a,d) to RH
15  rules trans_rs: # rule set defining transitive closure
16    path(x,y) if edge(x,y) # path holds for (x,y) if edge holds for (x,y)
17    path(x,y) if edge(x,z), path(z,y) # ... if edge(x,z) holds and path(z,y) holds
18  def transRH(): # to return transitive RH and reflexive role pairs
19    return infer(path, edge=RH, rules=trans_rs) + {(r,r): r in ROLES}
20  def AuthorizedUsers(role): # to return users having a role transitively
21    return {u: u in USERS, r in ROLES | (u,r) in UR and (r,role) in transRH()}
22  ...
22 h = new(HierRBAC, []) # create HierRBAC object h, with no args to setup
23 h.AddRole('chair') # call AddRole of h with role 'chair'
24  ...
24 h.AuthorizedUsers('chair') # call AuthorizedUsers of h with role 'chair'
25  ...

```

Fig. 1. An example program in Alda, for role-based access control (RBAC), demonstrating logic rules used with sets, functions, updates, and objects.

There has been a significant amount of related research, as discussed in Section 5. Our work contains two main contributions:

- A language that supports direct use of logic rules with sets, functions, updates, and objects, all as built-ins, seamlessly integrated, with a formal semantics.
- A compilation framework for implementation in a widely used programming language, where additional optimizations for rules can be exploited when available.

We have developed a prototype implementation of the compilation framework for Alda and experimented with a variety of programming and performance benchmarks. Our experiments strongly confirm the power and benefit of a seamlessly integrated language and the generally good performance of the implementation. Our implementation and benchmarks are publicly available (Tong et al. 2023).

2 Alda language

We first introduce rules and then describe how our overall language supports rules with sets and functions as well as imperative updates and object-oriented programming. Figure 1 shows an example program in Alda that uses all of rules, sets, functions, updates, and objects. It will be explained throughout Sections 2.1–2.6 when used as examples. A complete exposition of the formal semantics is in (Liu et al. 2023, Appendix A).

2.1 Logic rules

We support rule sets of the following form, where *name* is the name of the rule set, *declarations* is a set of predicate declarations, and the body is a set of rules.

```

rules name (declarations):
  rule+

```

A *rule* is either one of the two equivalent forms below (for users accustomed to either form), meaning that if *hypothesis*₁ through *hypothesis*_{*h*} all hold, then *conclusion* holds.

conclusion **if** *hypothesis*₁, *hypothesis*₂, ..., *hypothesis*_{*h*}
if *hypothesis*₁, *hypothesis*₂, ..., *hypothesis*_{*h*}: *conclusion*

If a conclusion holds without a hypothesis, then **if** and **:** are omitted.

Declarations are about predicates used in the rule set, for advanced uses, and are optional. For example, they may specify argument types of predicates, so rules can be compiled to efficient standalone imperative programs (Liu and Stoller 2009) that are expressed in typed languages (Rothamel and Liu 2007). They may also specify assumptions about predicates (Liu and Stoller 2020) to support different desired semantics (Liu and Stoller 2021; 2022). We omit the details because they are orthogonal to the focus of the paper. In particular, we omit types to avoid unnecessary clutter in code.

We use Datalog rules (Abiteboul *et al.* 1995; Maier *et al.* 2018) in examples, but our method of integrating semantics applies to rules in general. Each hypothesis and conclusion in a rule is an *assertion*, of the form

$$p(\mathit{arg}_1, \dots, \mathit{arg}_a),$$

where *p* is a *predicate*, and each *arg*_{*k*} is a variable or a constant. We use numbers and quoted strings to represent constants, and the rest are variables. As is standard for safe rules, all variables in the conclusion must be in a hypothesis. If a conclusion holds without a hypothesis, then each argument in the conclusion must be a constant, in which case the conclusion is called a *fact*. Note that a predicate is also called a *relation*, relating the arguments of the predicate.

Example. For computing the transitive closure of a graph in the running example, the rule set, named `trans_rs`, in Figure 1 (lines 15–17) can be written. The rules are the same as in dominant logic languages except for the use of lower-case variable names, the change of `:-` to **if**, and the omission of dot at the end of each rule. ■

Terminology. Consider a set of rules. Predicates not in any conclusion are called *base predicates*, and the other predicates are called *derived predicates*. We say that a predicate *p* *depends on* a predicate *q* if *p* is in the conclusion of a rule whose hypotheses contain *q* or contain a predicate that depends on *q* recursively. We say that a derived predicate *p* *fully depends on* a set *s* of base predicates if *p* does not depend on other base predicates.

Example. In rule set `trans_rs`, `edge` is a base predicate, and `path` is a derived predicate. `path` depends on `edge` and itself. `path` fully depends on `edge`. ■

2.2 Integrating rules with sets, functions, updates, and objects

Our overall language supports all of rule sets and the following language constructs as built-ins; all of them can appear in any scope—global, class, and local.

- Sets and set expressions (comprehension, aggregation, quantification, and high-level operations such as union) to make non-recursive queries over sets easy to express.
- Function and procedure definitions with optional keyword arguments, and function and procedure calls.
- Imperative updates by assignments and membership changes, to sets and data of other types, in sequencing, branching, and looping statements.

- Class definitions containing object field and method (function and procedure) definitions, object creations, and inheritance.

A name holding any value is *global* if it is introduced (declared or defined) at the global scope; is an *object field* if it is introduced for that object; or is *local* to the function, method, or rule set that contains it otherwise. After a name is defined, the value that it is holding is available: globally for a global name, on the object for an object field, and in the enclosing function, method, or rule set for a local name.

Example. Rule set `trans.rs` in Figure 1 (defined on lines 15–17 and queried using a call to an inference function. `infer`, on line 19) is used together with sets (defined on lines 3 and 12), set expressions (on lines 8, 19, and 21), functions (defined on lines 7–9, 18–19, and 20–21), procedures (defined on lines 2–3, 5–6, 10–12, and 13–14), updates (on lines 3, 6, 12, 14), classes (defined on lines 1 and 9, with inheritance), and objects (created on line 22). No extra code is needed to convert `edge` and `path`, declare logic variables, and so on. ■

The key ideas of our seamless integration of rules with sets, functions, updates, and objects are: (1) a predicate is a set-valued variable that holds the set of tuples for which the predicate is true, (2) queries using rules are calls to an inference function that computes desired sets using given sets, (3) values of predicates can be updated either directly as for other variables or by the inference function, and (4) predicates and rule sets can be object attributes as well as global and local names, just as sets and functions can.

Integrated semantics, ensuring declarative semantics of rules. In our overall language, the meaning of a rule set `rs` is completely declarative, exactly following the standard least fixed-point semantics of rules (Fitting 2002; Liu and Stoller 2009):

Given values of any set `s` of base predicates in `rs`, the meaning of `rs` is, for all derived predicates in `rs` that fully depend on `s`, the least set of values that can be inferred, directly or indirectly, by using the given values and the rules in `rs`;

for any derived predicate in `rs` that does not fully depend on `s`, that is, depends on any base predicate whose values are not given, its value is *undefined*.

The operational semantics for the rest of the language ensures this declarative semantics of rules. The precise constructs for using rules with sets, functions, updates, and objects are described in Sections 2.3–2.6.

2.3 Predicates as set-valued variables

For rules to be easily used with everything else, our most basic principle in designing the language is to treat a predicate as a set-valued variable that holds the set of tuples that are true for the predicate, that is:

For any predicate `p` over values `x1, ..., xa`, assertion `p(x1, ..., xa)` is true—that is, `p(x1, ..., xa)` is a fact—if and only if tuple `(x1, ..., xa)` is in set `p`. Formally,

$$p(x_1, \dots, x_a) \iff (x_1, \dots, x_a) \text{ in } p$$

This means that, as variables, predicates in a rule set can be introduced in any scope— as global variables, object fields, or variables local to the rule set—and they can be written into and read from without needing any extra interface.

Example. In rule set `trans.rs` in Figure 1, predicate `edge` is exactly a variable holding a set of pairs, such that `edge(x,y)` is true iff `(x,y)` is in `edge`, and `edge` is local to `trans.rs`.

In general, `edge` can be a global variable, an object field, or a local variable of `trans.rs`. Similarly for predicate `path`. ■

Writing to predicates is discussed later under updates to predicates, but reading and using values of predicates can simply use all operations on sets. We use set expressions including the following:

<code>exp in sexp</code>	membership
<code>exp not in sexp</code>	negated membership
<code>sexp₁ + sexp₂</code>	union
<code>{exp: v₁ in sexp₁, ..., v_k in sexp_k bexp}</code>	comprehension
<code>agg sexp</code> , where <i>agg</i> is count, max, min, sum	aggregation
<code>some v₁ in sexp₁, ..., v_k in sexp_k bexp</code>	existential quantification

A comprehension returns the set of values of `exp` for all combinations of values of variables that satisfy all membership clauses `vi in sexpi` and condition `bexp`. An aggregation returns the count, max, etc. of the set value of `sexp`. An existential quantification returns true iff for some combination of values of variables that satisfies all `vi in sexp` clauses, condition `bexp` holds. When an existential quantification returns true, variables `v1, ..., vk` are bound to a witness. Note that these set queries, as in (Liu *et al.* 2017), are more powerful than those in Python.

Example. For computing the transitive closure `T` of a set `E` of edges, the following `while` loop with quantification can be used (we will see that we use objects and updates as in Python except for the syntax `:=` for assignment in this paper):

```
T := E.copy()
while some (x,z) in T, (z,y) in E | (x,y) not in T:
    T.add((x,y))
```

In the comprehension and aggregation forms, each `vi` can also be a tuple pattern that elements of the set value of `sexpi` must match (Liu *et al.* 2017). A *tuple pattern* is a tuple in which each component is a non-variable expression, a variable possibly prefixed with `=`, a wildcard `_`, or recursively a tuple pattern. For a value to match a tuple pattern, it must have the corresponding tuple structure, with corresponding components equal the values of non-variable expressions and variables prefixed with `=`, and with corresponding components assigned to variables not prefixed with `=`; multiple occurrences of a variable must be assigned the same value; corresponding components of wildcard are ignored.

Example. To return the set of second component of pairs in `path` whose first component equals the value of variable `x`, and where that second component is also the first component of pairs in `edge` whose second component is 1, one may use a set comprehension with tuple patterns:

```
{y: (=x,y) in path, (y,1) in edge}
```

Now that predicates in rules correspond to set-valued variables, instead of functions or procedures, we can further see that logic variables, that is, variables in arguments of predicates in rules, are like pattern variables, that is, variables not prefixed with `=` in patterns. These variables are used for relating values, through what is generally called unification; they do not hold values, unlike variables prefixed with `=` in patterns.

2.4 Queries as calls to an inference function

For inference and queries using rules, calls to a built-in inference function `infer`, of the following form, are used, with `queryk`'s and `pk=sexpk`'s being optional:

```
infer(query1, ..., queryj, p1=sexp1, ..., pi=sexpi, rules=rs)
```

rs is the name of a rule set. Each $sexp_k$ is a set-valued expression. Each p_k is a base predicate of rs and is local to rs . Each $query_k$ is of the form $p(arg_1, \dots, arg_a)$, where p is a derived predicate of rs , and each argument arg_k is a constant, a variable possibly prefixed with $=$, or wildcard $_$. A variable prefixed with $=$ indicates a bound variable whose value will be used as a constant when evaluating the query. So arguments of queries are patterns too. If all arg_k 's are $_$, the abbreviated form p can be used.

Function `infer` can be called implicitly by the language implementation or explicitly by the user. It is called automatically as needed and can be called explicitly when desired.

Example. For inference using rule set `trans_rs` in Figure 1, where `edge` and `path` are local variables, `infer` can be called in many ways, including:

```
infer(path, edge=RH, rules=trans_rs)
infer(path(_,_), edge=RH, rules=trans_rs)
infer(path(1,_), path(_,=R), edge=RH, rules=trans_rs)
```

The first is as in Figure 1 (line 19). The first two calls are equivalent: `path` and `path(_,_)` both query the set of pairs of vertices having a path from the first vertex to the second vertex, following edges given by the value of variable `RH`. In the third call, `path(1,_)` queries the set of vertices having a path from vertex 1, and `path(_,=R)` queries the set of vertices having a path to the vertex that is the value of variable `R`.

If `edge` or `path` is a global variable or an object field, one may call `infer` on `trans_rs` without assigning to `edge` or querying `path`, respectively. ■

The operational semantics of a call to `infer` is exactly like other function calls, except for the special forms of arguments and return values, and of course the inference function performed inside:

- (1) For each value k from 1 to i , assign the set value of expression $sexp_k$ to predicate p_k that is a base predicate of rule set rs .
- (2) Perform inference using the rules in rs and the given values of base predicates of rs following the declarative semantics, including assigning to derived predicates that are not local.
- (3) For each value k from 1 to j , return the result of query $query_k$ as the k th component of the return value. The result of a query with l distinct variables not prefixed with $=$ is a set of tuples of l components, one for each of the distinct variables in their order of first occurrence in the query.

Note that when there are no $p_{k=sexp_k}$'s, only defined values of base predicates that are not local to rs are used; and when there are no $query_k$'s, only values of derived predicates that are not local to rs may be inferred and no value is returned. This is the case for implicit calls to `infer` on rs .

2.5 Updates to predicates

Values of base predicates can be updated directly as for other set-valued variables, and values of derived predicates are updated by the inference function.

Base predicates of a rule set rs that are local to rs are assigned values at calls to `infer` on rs , as described earlier. Base predicates that are not local can be updated by assignment statements or set update operations. We use

```
lexp := exp
```

for assignments, where ιexp can also be a nested tuple of variables, and each variable is assigned the corresponding component of the value of exp .

Derived predicates of a rule set rs can be updated only by calls to the inference function on rs . The updates must ensure the declarative semantics of rs :

Whenever a base predicate of rs is updated in the program, the values of the derived predicates in rs are maintained according to the declarative semantics of rs by calling `infer` on rs .

Updates to derived predicates of rs outside rs are not allowed, and any violation will be detected and reported at compile time if possible and at runtime otherwise.

Simply put, updates to base predicates trigger updates to derived predicates, and other updates to derived predicates are not allowed. This ensures the invariants that the derived predicates hold the values defined by the rule set based on values of the base predicates, as required by the declarative semantics. Note that this is the most straightforward semantics, but the implementation can avoid many inefficiencies with optimizations.

Example. Consider rule set `trans_rs` in Figure 1. If `edge` is not local, one may assign a set of pairs to `edge`:

```
edge := {(1,8), (2,9), (1,2)}
```

If `edge` is local, the calls to `infer` in the example in Section 2.4 assign the value of `RH` to `edge`.

If `path` is not local, then a call `infer(edge=RH, rules=trans_rs)` updates `path`, contrasting the first two calls to `infer` in the example in Section 2.4 that return the value of `path`.

If `path` is local, the return value of `infer` can be assigned to variables. For example, for the third call to `infer` in the example in Section 2.4, this can be

```
from1,toR := infer(path(1,_), path(_)=R, edge=RH, rules=trans_rs)
```

If both `edge` and `path` are not local, then whenever `edge` is updated, an implicit call `infer(rules=trans_rs)` is made automatically to update `path`. ■

For the RBAC example in Figure 1, different ways of using rules are possible, including (1) `alloc`: adding a rule `path(x,x) if role(x,x)` to the rule set, adding `role=ROLES` in the call to `infer`, and removing the union in function `transRH`, so all predicates are local variables; (2) `nonloc`: as in `alloc`, except to replace predicates `edge`, `role`, and `path` with `RH`, `ROLES`, and a new field `transRH`, respectively, replace call `transRH()` with field `transRH`, and remove function `transRH`; (3) `union`: as in Figure 1; and other combinations of aspects of (1)–(3).

2.6 Using predicates and rules with objects and classes

Predicates and rule sets can be object fields as well as global and local names, just as sets and functions can, as discussed in Section 2.2. This allows predicates and rule sets to be used seamlessly with objects in object-oriented programming.

For other constructs than those described above, we use those in high-level object-oriented languages. We mostly use Python syntax (looping, branching, indentation for scoping, ‘:’ for elaboration, ‘#’ for comments, etc.) for succinctness, but with a few conventions from Java (keyword `new` for object creation, keyword `extends` for subclassing, and omission of `self`, the equivalent of `this` in Java, when there is no ambiguity) for ease of reading.

Example. We use role-based access control (RBAC) to show the need of using rules with all of sets, functions, updates, and objects and classes.

RBAC is a security policy framework for controlling user access to resources based on roles and is widely used in large organizations. The ANSI standard for RBAC (ANSI INCITS 2004) was approved in 2004 after several rounds of public review (Sandhu et al. 2000; Jaeger and Tidswell 2000; Ferraiolo et al. 2001), building on much research during the preceding decade and earlier. High-level executable specifications were developed for the entire RBAC standard (Liu and Stoller 2007), where all queries are declarative except for computing the transitive role-hierarchy relation in Hierarchical RBAC, which extends Core RBAC.

Core RBAC defines functionalities relating users, roles, permissions, and sessions. It includes the sets and update and query functions in class `CoreRBAC` in Figure 1, as in (Liu and Stoller 2007).¹

Hierarchical RBAC adds support for a role hierarchy, `RH`, and update and query functions extended for `RH`. It includes the update and query functions in class `HierRBAC` in Figure 1, as in (Liu and Stoller 2007),¹ except that function `transRH()` in (Liu and Stoller 2007) computes the transitive closure of `RH` plus reflexive role pairs for all roles in `ROLES` by using a complex and inefficient `while` loop much worse than that in Section 2.3 (due to Python's lack of `some` with witness) plus a union with the set of reflexive role pairs $\{(r,r): r \text{ in } \text{ROLES}\}$, whereas function `transRH()` in Figure 1 simply calls `infer` and unions the result with reflexive role pairs.

Note though, in the RBAC standard, a relation `transRH` is used in place of `transRH()`, intending to maintain the transitive role hierarchy incrementally while `RH` and `ROLES` change. It is believed that this is done for efficiency, because the result of `transRH()` is used continually, while `RH` and `ROLES` change infrequently. However, the maintenance was done inappropriately (Liu and Stoller 2007; Li et al. 2007) and warranted the use of `transRH()` to ensure correctness before efficiency.

Overall, the RBAC specification relies extensively on all of updates, sets, functions, and objects and classes with inheritance, besides rules: (1) updates for setting up and updating the state of the RBAC system, (2) sets and set expressions for holding the system state and expressing set queries exactly as specified in the RBAC standard, (3) methods and functions for defining and invoking update and query operations, and (4) objects and classes for capturing different components—`CoreRBAC`, `HierRBAC`, constraint RBAC, their further refinement, extensions, and combinations, totaling 9 components, corresponding to 9 classes, including 5 subclasses of `HierRBAC` (ANSI INCITS 2004; Liu and Stoller 2007). ■

3 Compilation

We describe our compilation framework for implementing Alda, by building on an object-oriented language that supports all features except rules and queries and on an efficient logic rule engine for queries using rules. Three main tasks are (1) compiling rule sets to generate rules accepted by the rule engine, (2) compiling queries using rules to generate queries accepted by the rule engine, together with automatic conversion of data and query results, and (3) compiling updates to predicates that require implicit automatic queries and updates of the query results. The compiler must appropriately handle scoping of rule

¹ Only a few selected sets and functions are included, and with small changes to names and syntax.

sets and predicates for all three tasks. Besides that, task (1) is straightforward, task (2) is also straightforward but tedious, and task (3) requires the most analysis, so we focus on task (3) below.

We first describe how to compile all possible updates to predicates, starting with the checks and actions needed to correctly handle updates for a single rule set with implicit and explicit calls to `infer`. We then describe how to implement the inference in `infer`. In (Liu *et al.* 2023, Appendix B), we systematize powerful optimizations that can be added in the overall compilation framework; clearly separated handling of updates and queries in our compilation framework allows optimizations to be added in a modular fashion.

3.1 Compiling updates to predicates

The operational semantics to ensure the declarative semantics of a rule set rs is conceptually simple, but for efficiency, the implementation required varies, depending on the kind of updates to base predicates of rs outside rs . Note that inside rs there are no updates to base predicates of rs , by definition of base predicate.

- (1) **Local updates.** Local variables of rs , that is, predicates local to rs , can be assigned values only at explicit calls to `infer` on rs . Such a call passes in values of local variables that are base predicates of rs before doing the inference. Values of local variables that are derived predicates of rs can only be used in constructing answers to the queries in the call, and the answers are returned from the call.

There are no updates outside rs to local variables that are derived predicates of rs , by definition of local variables.

- (2) **Non-local updates.** For updates to non-local variables of rs , an implicit call to `infer` on rs needs to be made only after every update to a base predicate of rs . Statements outside rs that update derived predicates of rs are identified and reported as errors.

In languages or application programs where variables hold data values, such as in database languages and applications, these updates can be determined simply at compile time, for example, if s holds a set value, then $s := s + \{x\}$ updates the set value of s . This is also the case when logic rules are used in these languages and programs.

In programs where variables may be references to data values, each update needs to check whether the updated variable may alias a predicate of rs , conservatively at compile-time if possible, and at runtime otherwise.

To satisfy these requirements, the overall method for compiling an update to a variable v outside rule sets is:

- In languages or application programs where variables hold data values, report a compile-time error if v is a derived predicate of any rule set; otherwise, for each rule set rs that contains v as a base predicate, insert code, after the update, that calls `infer` on rs with no arguments for base predicates and no queries.
- Otherwise, if v may refer to a predicate in a rule set, insert code that does the following after the update: if v refers to a derived predicate of any rule set, report a runtime error and exit; otherwise for each rule set rs , if v refers to a base predicate of rs , call `infer` on rs with no arguments for base predicates and no queries.

Our method for compiling an explicit call to `infer` on a rule set directly follows the operational semantics of `infer`.

In effect, function `infer` is called to implement a wide range of control: from inferring everything possible using all rule sets and values of all base predicates at every update, to answering specific queries using specific rules and specific sets of values of specific base predicates at explicit calls.

Obviously, updates in different cases may have significant impact on program efficiency. Update analysis is needed to determine the case and generate correct code. Our compilation method above minimizes calls to `infer` in each case.

3.2 Implementing inference and queries

Any existing method can be used to implement the functionality inside `infer`. The inference and queries for a rule set can use either bottom-up or top-down evaluation (Kifer and Liu 2018; Tekle and Liu 2010; 2011), so long as they use the rule set and values of the base predicates according to the declarative semantics of rules.

The inference and queries can be either performed by using a general logic rule engine, for example, XSB (Sagonas et al. 1994; Swift et al. 2022), or compiled to specialized standalone executable code as in, for example, (Liu and Stoller 2009; Rothamel and Liu 2007; Jordan et al. 2016), that is then executed. Our current implementation uses the former approach, by indeed using the well-known XSB system, as described in Section 4, because it allows easier extensions to support more kinds of rules and optimizations that are already supported in XSB. Other powerful logic rule engines, including efficient answer set programming (ASP) systems such as Clingo (Gebser et al. 2019), can certainly be used also.

4 Implementation and experimental evaluation

We have implemented a prototype compiler for Alda. The compiler generates executable code in Python. The generated code calls the XSB logic rule engine (Sagonas et al. 1994; Swift et al. 2022) for inference using rules.

We implemented Alda by extending the DistAlgo compiler (Liu et al. 2012; 2017; Lin and Liu 2022). DistAlgo is an extension of Python with high-level set queries as well as distributed processes. The compiler is implemented in Python 3, and uses the Python parser. So Python syntax is used in place of the ideal syntax presented in Section 2, allowing any user with Python to run Alda directly.

The Alda implementation extends the DistAlgo compiler to support rule-set definitions, function `infer`, and maintenance of derived predicates at updates to non-local variables. It handles direct updates to variables used as predicates, not updates through aliasing, as we found this to be the only update case in all benchmarks and other examples we have seen; we think this is because using logic rules with updates is similar to using queries and updates in relational databases, with no need of updates through aliasing. Currently Datalog rules extended with unrestricted negation are supported, and well-founded semantics computed by XSB is used; extensions for more general rules can be handled similarly, and inference using XSB can remain the same. Calls to `infer` are automatically added at updates to non-local base predicates of rule sets.

In particular, the following Python syntax is used for rule sets, where a rule can be either one of the two forms below, so the only restriction is that the name `rules` is reserved.

```
def rules (name = rsname):  
    conclusion, if_(hypothesis1, hypothesis2, ..., hypothesisn)  
    if (hypothesis1, hypothesis2, ..., hypothesisn): conclusion  
    ...
```

Rule sets are translated into Prolog rules at compile time. The directive `:- auto_table.` is added for automatic tabling in XSB.

For function `infer`, the implementation translates the values of predicates and the list of queries into facts and queries in standard Prolog syntax, and translates the query answers back to values of set variables. It invokes XSB using a command line in between, passing data through files; this external interface has an obvious overhead, but it has not affected Alda having generally good performance. `infer` automatically reads and writes non-local predicates used in a rule set.

Note that the overhead of the external interface can be removed with an in-memory interface from Python to XSB, which is actively being developed by the XSB team.² However, even with the overhead of the external interface, Alda is still faster or even drastically faster than half or more of the rule engines tested in OpenRuleBench (Liang *et al.* 2009) for all benchmarks measured except DBLP (even though OpenRuleBench uses the fastest manually optimized program for each problem for each rule engine), and than not using rules at all (without manually writing or adapting a drastically more complex, specialized algorithm implementation for each problem).

Building on top of DistAlgo and XSB, the compiler consists of about 1100 lines of Python and about 50 lines of XSB. This is owing critically to the overall framework and comprehensive support, especially for high-level queries, already in the DistAlgo compiler and to the powerful query engine of XSB. The parser for the rule extension is about 270 lines, and update analysis and code generation for rules and inference are about 800 lines.

The current compiler does not perform further optimizations, because they are orthogonal to the focus of this paper, and our experiments already showed generally good performance. Further optimizations can be implemented in either the Alda compiler to generate optimized rules and tabling and indexing directives, or in XSB. Incremental maintenance under updates can also be implemented in either one, with a slightly richer interface between the two.

We discuss our experiments on the benchmarks summarized in Table 1. Detailed description of the benchmarks are in (Liu *et al.* 2022; 2023). Just as the benchmarks selected, the experiments selected are also meant to show generally good performance even under the most extreme overhead penalties we have encountered—runs with large data (DBLP and PA), large query results (transitive closure TC), large rules (Wine), frequent switches among different ways of using rules and other features (RBAC and PA), and frequent external invocations of the rule engine (RBAC). Our extensive experiments with other uses of Alda have experienced minimum performance overhead.

² A version for Unix, not yet Windows, has been released: passing data of size 100 million in memory took about 30 nanoseconds per element (Swift *et al.* 2022, release notes). So even the largest data in our experiments, of size a few millions, would take 0.1–0.2 seconds to pass in memory, instead of 10–20 seconds with the current external interface.

Table 1. *Benchmarks from different kinds of problems. RBAC benchmarks are for different ways of using rules as at the end of Section 2.4. PA is a mixture of problems from class hierarchy analysis. Under Variants, suffixes py and da indicate using while loops like that in Section 2.3 in Python and DistAlgo, respectively, instead of using rules.*

Benchmark sets	Benchmarks	Variants and timing	Problem kinds	Code/data size
Open-RuleBench (Liang et al. 2009)	13 incl. LUBM, Mondial, DBLP, TC, WordNet, Wine	TCrev, TCda, TCpy, ORBtimer	Many kinds of rules and queries, but missing aggregate queries	Largest rule set: 967 rules, Largest data size: 2.4M+
RBAC as in Section 2.6	RBACalloc, RBACnonloc, RBACunion	RBACda, RBACpy, RBACtimer	Interleaved object queries and updates with function and recursive rules	Program size: 385–423, Data size: 10K+
Program Analysis	PA (on any prog.: numpy, pandas, matplotlib, pytorch, sympy, etc.)	PAopt, PATimer	Interleaved rules, aggregate and set queries, and recursive functions	Program size: 55 XSB, 33 Alda, largest data size: 5.1M+

All measurements were taken on a machine with an Intel Xeon X5690 3.47 GHz CPU, 94 GB RAM, running 64-bit Ubuntu 16.04.7, Python 3.9.9, and XSB 4.0.0. For each experiment, the reported running times are CPU times averaged over 10 runs. Garbage collection in Python was disabled for smoother running times when calling XSB. Program sizes are numbers of lines excluding comments and empty lines. Data sizes are number of facts.

We summarize the results from the experiments below. Detailed measurements and explanations are in (Liu et al. 2022; 2023).

- Compared with XSB programs in OpenRuleBench, the corresponding Alda programs are much smaller, almost all by dozens or even hundreds of lines, because all benchmarking code is in a single shared 45-line ORBtimer, much easier in Python than XSB. Compilation times are all 0.6 seconds or less.
- Running times for all benchmarks and variants, except for PA, are as expected, for example, TC is drastically faster than TCpy and TCda, and essentially as fast as XSB if not for the overhead of using external interface with XSB; and RBACnonloc is much faster than RBACalloc due to updates being much less frequent than queries. The overhead of using external interface is obvious: for example, for TC, up to 5.9 seconds, out of 29.2 seconds, for graphs of 100K edges; for PA, 13.1 seconds, out of 15.2 seconds, on the largest program, SymPy; and worst for DBLP, 26.9 seconds, out of 30.6 seconds, on over 2.4M facts.

However, even so, Alda is competitive, as described above, and the overhead is expected to be reduced to 1% of it with an in-memory Python–XSB interface.

- For PA, the corresponding XSB programs were all slower and even drastically slower than Alda programs, even 120 times slower on PyTorch. Significant effort was spent

on performance debugging and manual optimization before we eventually created a version that is faster than Alda—5.1 vs. 15.2 seconds on SymPy.

5 Related work and conclusion

There has been extensive effort in design and implementation of languages to support programming with logic rules together with other programming paradigms, by extending logic languages, extending languages in other paradigms, or developing multi-paradigm or other standalone languages.

A large variety of logic rule languages have been extended to support sets, functions, updates, and/or objects, etc. (Kifer and Liu 2018; Körner *et al.* 2022). For example, see Maier *et al.* (2018) for Datalog and variants extended with sets, functions, objects, updates, higher-order extensions, and more. In particular, many Prolog variants support sets, functions, updates, objects, constraints, etc. For example, Prolog supports `assert` for updates, as well as `cut` and `negation as failure` that are imperative instead of declarative (Sterling and Shapiro 1994); Flora (Yang and Kifer 2000; Kifer *et al.* 2020) builds on XSB and supports objects (F-logic), higher-order programming (HiLog), and updates (Transaction Logic); and Picat (Zhou 2016) builds on B-Prolog and supports updates, comprehensions, etc. Lambda Prolog (Miller and Nadathur 2012) extends Prolog with simply typed lambda terms and higher-order programming. Functional logic languages, such as Mercury (Somogyi *et al.* 1995) and Curry (Hanus 2013), combine functional programming and logic programming. Some logic programming systems are driven by scripting externally, for example, using Lua for IDP (Bruynooghe *et al.* 2014), and shell scripts for LogicBlox (Aref *et al.* 2015). Additional examples of Datalog extensions include Flix (Madsen *et al.* 2016; Madsen and Lhoták 2020), which supports lattices and monotone functions, and DDlog (Ryzhyk and Budiú 2019), which supports incremental maintenance under updates to input relations. These languages and extensions do not support predicates as set-valued variables together with commonly used updates and objects in a simple and direct way, or do not support them at all.

Many languages in other programming paradigms, especially including imperative languages and object-oriented languages, have been extended to support rules by being a host language. This is generally through explicit library interfaces of the host languages to connect with a particular logic language, for example, a Java interface for XSB through InterProlog (Calejo 2004; Swift *et al.* 2022), C++ and Python interfaces for ASP systems dlhex (Redl 2016) and Potassco (Banbara *et al.* 2017), a Python interface for IDP (Vennekens 2017), Rust and other interfaces for DDlog (Ryzhyk and Budiú 2019), and many more, for example, for miniKanren (Byrd 2009). Hosting logic languages through explicit interfaces requires programmers to write extra wrapper code for going to the rule language and coming back—declare predicates and/or logic variables, wrap features in special objects, functions, macros, etc., and/or convert data to and from special representations. They are in the same spirit as interfaces such as JDBC (Reese 2000) for using database systems from languages such as Java.

Multi-paradigm languages and other standalone languages have also been developed. For example, the Mozart system for the Oz multi-paradigm programming language (Roy and Haridi 2004) supports logic, functional, and constraint as well as imperative and concurrent programming. However, it is similar to logic languages extended with other features, because it supports logic variables, but not state variables to be assigned

to as in commonly used imperative languages. Examples of other languages involving logic and constraints with updates and/or objects include LOGRES (Cacace et al. 1990), which integrates object-oriented data modeling and updates with rules under inflationary semantics; TLA+ (Lamport 1994), a logic language for specifying actions; CLAIRE (Caseau et al. 2002), an object-oriented language that supports functions, sets, and rules whose conclusions are actions; LINQ (Meijer et al. 2006; LINQ 2023), an extension of C# for SQL-like queries; IceDust (Harkes et al. 2016), a Java-based language for querying data with path-based navigation and incremental computation; extended LogiQL in SolverBlox (Borraz-Sánchez et al. 2018), for mathematical and logic programming on top of Datalog with updates and constraints; and other logic-based query languages, for example, Datomic (Anderson et al. 2016) and SOUL (De Roover et al. 2011). These are either logic languages lacking general imperative and object-oriented programming constructs, or imperative and object-oriented languages lacking the power and full declarativeness of logic rules.

In conclusion, Alda supports ease of programming with logic rules together with all of sets, functions, updates, and objects as seamlessly integrated built-ins, without extra interfaces or boiler-plate code. As a direction for future work, many optimizations can be added to improve the efficiency of implementations. This includes optimizing the logic rule engines used (Liu and Stoller 2009; Tekle and Liu 2011), the interfaces and interactions with them, and using other efficient rule systems such as Clingo (Gebser et al. 2019) and specialized rule implementations such as Souffle (Jordan et al. 2016) to obtain the best possible performance.

Acknowledgments

We thank David S. Warren for an initial 28-line XSB program for interface to XSB, and Tuncay Tekle for help implementing some benchmarks and running some preliminary experiments. We also thank Thang Bui for additional applications in program analysis and optimization, and students in undergraduate and graduate courses for using Alda and its earlier versions, called DA-rules.

This work was supported in part by NSF under grants CCF-1954837, CCF-1414078, and IIS-1447549 and ONR under grants N00014-21-1-2719, N00014-20-1-2751, and N00014-15-1-2208.

References

- ABITEBOUL, S., HULL, R. AND VIANU, V. 1995. *Foundations of Databases: The Logical Level*. Addison-Wesley.
- ANDERSON, J., GAARE, M., HOLGUÍN, J., BAILEY, N. AND PRATLEY, T. 2016. The Datomic database. In *Professional Clojure*. Wiley Online Library, Chapter 6, 169–215.
- ANSI INCITS. 2004. Role-Based Access Control. ANSI INCITS 359-2004, American National Standards Institute, International Committee for Information Technology Standards.
- AREF, M., TEN CATE, B., GREEN, T. J., KIMELFELD, B., OLTEANU, D., PASALIC, E., VELDUIZEN, T. L. AND WASHBURN, G. 2015. Design and implementation of the LogicBlox system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. 1371–1382. <https://doi.org/10.1145/2723372.2742796>.

- BANBARA, M., KAUFMANN, B., OSTROWSKI, M. AND SCHAUB, T. 2017. Clingcon: The next generation. *Theory and Practice of Logic Programming* 17, 4, 408–461.
- BORRAZ-SÁNCHEZ, C., KLABJAN, D., PASALIC, E. AND AREF, M. 2018. SolverBlox: Algebraic modeling in Datalog. In *Declarative Logic Programming: Theory, Systems, and Applications*, M. Kifer and Y. A. Liu, Eds. ACM and Morgan & Claypool, Chapter 6, 331–356. <https://doi.org/10.1145/3191315.3191322>.
- BRUYNNOOGHE, M., BLOCKEEL, H., BOGAERTS, B., DE CAT, B., DE POOTER, S., JANSEN, J., LABARRE, A., RAMON, J., DENECKER, M. AND VERWER, S. 2014. Predicate logic as a modeling language: Modeling and solving some machine learning and data mining problems with IDP3. *Theory and Practice of Logic Programming* 15, 6, 783–817. <https://doi.org/10.1017/S147106841400009X>.
- BYRD, W. E. 2009. Relational programming in miniKanren: Techniques, applications, and implementations. Ph.D. thesis, Indiana University.
- CACACE, F., CERI, S., CRESPI-REGHIZZI, S., TANCA, L. AND ZICARI, R. 1990. Integrating object-oriented data modelling with a rule-based programming paradigm. In *Proceedings of the 1990 ACM SIGMOD International Conference on Management of Data*, 225–236.
- CALEJO, M. 2004. InterProlog: Towards a declarative embedding of logic programming in Java. In *Proceedings of the 9th European Conference on Logics in Artificial Intelligence*. LNCS, vol. 3229. Springer, 714–717.
- CASEAU, Y., JOSSET, F.-X. AND LABURTHER, F. 2002. Claire: Combining sets, search and rules to better express algorithms. *Theory and Practice of Logic Programming* 2, 6, 769–805.
- DE ROOVER, C., NOGUERA, C., KELLENS, A. AND JONCKERS, V. 2011. The SOUL tool suite for querying programs in symbiosis with Eclipse. In *Proceedings of the 9th International Conference on Principles and Practice of Programming in Java*, 71–80.
- FERRAILOLO, D. F., SANDHU, R., GAVRILA, S., KUHN, D. R. AND CHANDRAMOULI, R. 2001. Proposed NIST standard for role-based access control. *ACM Transactions on Information and Systems Security* 4, 3, 224–274.
- FITTING, M. 2002. Fixpoint semantics for logic programming: A survey. *Theoretical Computer Science* 278, 1, 25–51.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2019. Multi-shot ASP solving with clingo. *Theory and Practice of Logic Programming* 19, 1, 27–82. <https://doi.org/10.1017/S1471068418000054>.
- GEIGER, K. 1995. *Inside ODBC*. Microsoft Press.
- HANUS, M. 2013. Functional logic programming: From theory to Curry. In *Programming Logics*. Springer, 123–168.
- HARKES, D. C., GROENEWEGEN, D. M. AND VISSER, E. 2016. IceDust: Incremental and eventual computation of derived values in persistent object graphs. In *30th European Conference on Object-Oriented Programming*. LIPIcs, vol. 56. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 11:1–11:26.
- JAEGER, T. AND TIDSWELL, J. 2000. Rebuttal to the NIST RBAC model proposal. In *Proceedings of the 5th ACM Workshop on Role Based Access Control*. 66.
- JORDAN, H., SCHOLZ, B. AND SUBOTIĆ, P. 2016. Soufflé: On synthesis of program analyzers. In *Proceedings of the International Conference on Computer Aided Verification*. Springer, 422–430.
- KIFER, M. AND LIU, Y. A., Eds. 2018. *Declarative Logic Programming: Theory, Systems, and Applications*. ACM and Morgan & Claypool.
- KIFER, M., YANG, G., WAN, H. AND ZHAO, C. 2020. *Ergo Lite (a.k.a. Flora-2): User's Manual Version 2.1*. Stony Brook University. <http://flora.sourceforge.net/>. Accessed May 25, 2023.
- KÖRNER, P., LEUSCHEL, M., BARBOSA, J. A., COSTA, V. S., DAHL, V., HERMENEGILDO, M. V., MORALES, J. F., WIELEMAKER, J., DIAZ, D., ABREU, S. AND CIATTO, G. 2022. Fifty years of Prolog and beyond. *Theory and Practice of Logic Programming* 22, 6, 776–858.

- LAMPART, L. 1994. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems* 16, 3, 872–923.
- LI, N., BYUN, J.-W. AND BERTINO, E. 2007. A critique of the ANSI standard on role-based access control. *IEEE Security and Privacy* 5, 6, 41–49.
- LIANG, S., FODOR, P., WAN, H. AND KIFER, M. 2009. OpenRuleBench: An analysis of the performance of rule engines. In *Proceedings of the 18th International Conference on World Wide Web*. ACM Press, 601–610.
- LIN, B. AND LIU, Y. A. 2014 (Latest update January 30, 2022). DistAlgo: A language for distributed algorithms. <http://github.com/DistAlgo>. Accessed May 25, 2023.
- LINQ 2023. Language Integrated Query (LINQ). <https://docs.microsoft.com/dotnet/csharp/linq>. Accessed May 25, 2023.
- LIU, Y. A. 2018. Logic programming applications: What are the abstractions and implementations? In *Declarative Logic Programming: Theory, Systems, and Applications*, M. Kifer and Y. A. Liu, Eds. ACM and Morgan & Claypool, Chapter 10, 519–557. Also <https://arxiv.org/abs/1802.07284>.
- LIU, Y. A. AND STOLLER, S. D. 2007. Role-based access control: A corrected and simplified specification. In *Department of Defense Sponsored Information Security Research: New Methods for Protecting Against Cyber Threats*. Wiley, 425–439.
- LIU, Y. A. AND STOLLER, S. D. 2009. From Datalog rules to efficient programs with time and space guarantees. *ACM Transactions on Programming Languages and Systems* 31, 6, 1–38. <https://doi.org/10.1145/1552309.1552311>.
- LIU, Y. A. AND STOLLER, S. D. 2020. Founded semantics and constraint semantics of logic rules. *Journal of Logic and Computation* 30, 8 (Dec.), 1609–1638. Also <http://arxiv.org/abs/1606.06269>.
- LIU, Y. A. AND STOLLER, S. D. 2021. Knowledge of uncertain worlds: Programming with logical constraints. *Journal of Logic and Computation* 31, 1 (Jan.), 193–212. Also <https://arxiv.org/abs/1910.10346>.
- LIU, Y. A. AND STOLLER, S. D. 2022. Recursive rules with aggregation: A simple unified semantics. *Journal of Logic and Computation* 32, 8 (Dec.), 1659–1693. Also <http://arxiv.org/abs/2007.13053>.
- LIU, Y. A., STOLLER, S. D. AND LIN, B. 2017. From clarity to efficiency for distributed algorithms. *ACM Transactions on Programming Languages and Systems* 39, 3 (May), 12:1–12:41. Also <http://arxiv.org/abs/1412.8461>.
- LIU, Y. A., STOLLER, S. D., LIN, B. AND GORBOVITSKI, M. 2012. From clarity to efficiency for distributed algorithms. In *Proceedings of the 27th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications*, 395–410.
- LIU, Y. A., STOLLER, S. D., TONG, Y. AND LIN, B. 2023. Integrating logic rules with everything else, seamlessly. *Computing Research Repository arXiv:2305.19202 [cs.PL]*.
- LIU, Y. A., STOLLER, S. D., TONG, Y., LIN, B. AND TEKLE, K. T. 2022. Programming with rules and everything else, seamlessly. *Computing Research Repository arXiv:2205.15204 [cs.PL]*. <http://arxiv.org/abs/2205.15204>.
- LIU, Y. A., STOLLER, S. D., TONG, Y. AND TEKLE, K. T. 2023. Benchmarking for integrating logic rules with everything else. In *Proceedings of the 39th International Conference on Logic Programming (Technical Communications)*. Open Publishing Association.
- MADSEN, M. AND LHOTÁK, O. 2020. Fixpoints for the masses: Programming with first-class Datalog constraints. *Proceedings of the ACM on Programming Languages* 4, OOPSLA, 1–28.
- MADSEN, M., YEE, M.-H. AND LHOTÁK, O. 2016. From Datalog to Flix: A declarative language for fixed points on lattices. *ACM SIGPLAN Notices* 51, 6, 194–208.
- MAIER, D., TEKLE, K. T., KIFER, M. AND WARREN, D. S. 2018. Datalog: Concepts, history and outlook. In *Declarative Logic Programming: Theory, Systems, and Applications*, M. Kifer and Y. A. Liu, Eds. ACM and Morgan & Claypool, Chapter 1, 3–120.

- MEIJER, E., BECKMAN, B. AND BIERMAN, G. 2006. LINQ: reconciling object, relations and XML in the .NET framework. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, 706–706.
- MILLER, D. AND NADATHUR, G. 2012. *Programming with Higher-Order Logic*. Cambridge University Press.
- REDL, C. 2016. The DLVHEX system for knowledge representation: Recent advances (system description). *Theory and Practice of Logic Programming* 16, 5–6, 866–883.
- REESE, G. 2000. *Database Programming with JDBC and JAVA*. O'Reilly Media, Inc.
- ROTHAMEL, T. AND LIU, Y. A. 2007. Efficient implementation of tuple pattern based retrieval. In *Proceedings of the ACM SIGPLAN 2007 Workshop on Partial Evaluation and Program Manipulation*, 81–90. <https://doi.org/10.1145/1244381.1244394>.
- ROY, P. V. AND HARIDI, S. 2004. *Concepts, Techniques, and Models of Computer Programming*. MIT Press.
- RYZHYK, L. AND BUDIU, M. 2019. Differential datalog. In *Datalog 2.0, 3rd International Workshop on the Resurgence of Datalog in Academia and Industry*, 56–67.
- SAGONAS, K., SWIFT, T. AND WARREN, D. S. 1994. XSB as an efficient deductive database engine. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*. ACM Press, 442–453.
- SANDHU, R., FERRAILOLO, D. AND KUHN, R. 2000. The NIST model for role-based access control: Towards a unified standard. In *Proceedings of the 5th ACM Workshop on Role-Based Access Control*, 47–63.
- SOMOGYI, Z., HENDERSON, F. J. AND CONWAY, T. C. 1995. Mercury, an efficient purely declarative logic programming language. *Australian Computer Science Communications* 17, 499–512.
- STERLING, L. AND SHAPIRO, E. 1994. *The Art of Prolog*, 2nd ed. MIT Press.
- SWIFT, T., WARREN, D. S., SAGONAS, K., FREIRE, J., RAO, P., CUI, B., JOHNSON, E., DE CASTRO, L., MARQUES, R. F., SAHA, D., DAWSON, S. AND KIFER, M. 2022. *The XSB System Version 5.0.x*. <http://xsb.sourceforge.net>. Latest release May 12, 2022.
- TEKLE, K. T. AND LIU, Y. A. 2010. Precise complexity analysis for efficient Datalog queries. In *Proceedings of the 12th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming*, 35–44. <https://doi.org/10.1145/1836089.1836094>.
- TEKLE, K. T. AND LIU, Y. A. 2011. More efficient Datalog queries: Subsumptive tabling beats magic sets. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, 661–672. <http://doi.acm.org/10.1145/1989323.1989393>.
- TONG, Y., LIN, B., LIU, Y. A. AND STOLLER, S. D. 2023. ALDA. <http://github.com/DistAlgo/alda>. Accessed May 25, 2023.
- VENNEKENS, J. 2017. Lowering the learning curve for declarative programming: A Python API for the IDP system. In *Proceedings of 19th International Symposium on Practical Aspects of Declarative Languages*. Springer, 86–102.
- WARREN, D. S. AND LIU, Y. A. 2017. AppLP: A dialogue on applications of logic programming. *Computing Research Repository arXiv:1704.02375 [cs.PL]*.
- YANG, G. AND KIFER, M. 2000. FLORA: Implementing an efficient DOOD system using a tabling logic engine. In *Proceedings of the 1st International Conference on Computational Logic*. Springer, 1078–1093. https://doi.org/10.1007/3-540-44957-4_72.
- ZHOU, N.-F. 2016. Programming in Picat. In *Proceedings of the 10th International Symposium on Rule Technologies: Research, Tools, and Applications*. Springer, 3–18.