

Logic programming applications: What are the abstractions and implementations?*

Yanhong A. Liu

Computer Science Department, Stony Brook University

liu@cs.stonybrook.edu

December 24, 2017

Abstract

This article presents an overview of applications of logic programming, classifying them based on the abstractions and implementations of logic languages that support the applications. The three key abstractions are join, recursion, and constraint. Their essential implementations are for-loops, fixed points, and backtracking, respectively. The corresponding kinds of applications are database queries, inductive analysis, and combinatorial search, respectively. We also discuss language extensions and programming paradigms, summarize example application problems by application areas, and touch on example systems that support variants of the abstractions with different implementations.

1 Introduction

Common reasoning with logic is the root of logic programming, which allows logic rules and facts to be expressed formally and used precisely for inference, querying, and analysis in general. Logic formalisms, or languages, allow complex application problems to be expressed declaratively with high-level abstractions and allow desired solutions to be found automatically with potentially efficient low-level implementations.

The biggest challenge in logic programming has been the need for efficient implementations. Much progress has been made, with efficient implementations in some cases beating manually written low-level code. However, inadequate performance in many cases has led to

*This work was supported in part by NSF under grants CCF-0964196, CCF-1248184, CCF-1414078, and IIS-1447549; and ONR under grants N00014-15-1-2208.

the introduction of non-declarative features in logic languages and resulted in the writing of obscure logic programs.

Despite the challenges, the most exciting aspect of logic programming is its vast areas of applications. They range from database queries to program analysis, from text processing to decision making, from security to knowledge engineering, and more. These vast, complex, and interrelated areas make it challenging but necessary to provide a deeper understanding of the various kinds of applications in order to help advance the state of the art of logic programming and realize its benefits.

This article presents an overview of applications of logic programming based on a study of the abstractions and implementations of logic languages. The rationale is that abstractions and implementations are the enabling technologies of the applications. The abstractions are essential for determining what kinds of application problems can be expressed and how they can be expressed, for ease of understanding, reuse, and maintenance. The underlying implementations are essential for high-level declarative languages to be sufficiently efficient for substantial applications.

We discuss the following essential abstractions, where data abstractions are for expressing the data, and control abstractions are for expressing computations over the data:

1. data abstractions: objects and relationships;
2. control abstractions: (1) join, (2) recursion, and (3) constraint, which capture bounded, cyclic, and general computations, respectively.

In logic languages, the data abstractions as objects and relationships are essential for all three control abstractions.

The essential techniques for implementing the three control abstractions listed are (1) for-loops, (2) fixed points, and (3) backtracking, respectively. The corresponding kinds of applications are

- (1) database-style queries, e.g., for ontology management, business intelligence, and access control;
- (2) inductive analysis, e.g., for text processing, program analysis, network traversal, and trust management;
- (3) combinatorial search, e.g., for decision making, resource allocation, games and puzzles, and administrative policy analysis.

We categorize application problems using these three control abstractions because they capture conceptually different kinds of problems, with inherently different implementation techniques, and at the same time correspond to very different classes of applications.

Note that the same application domain may use different abstractions and implementations for different problems. For example, enterprise software may use all three of traditional

database queries, inductive analysis, and combinatorial search, for business intelligence and decision making; and security policy analysis and enforcement may use database-style queries for access control, inductive analysis for trust management, and combinatorial search for administrative policy analysis.

We also discuss additional extensions, especially regular-expression paths for higher-level queries and updates for modeling actions; additional applications; and abstractions used in main programming paradigms. We also touch on several well-known systems while discussing the applications.

There is a large body of prior work, including surveys of logic programming in general and applications in particular, as discussed in Section 7. This article distinguishes itself from past work by analyzing classes of applications based on the language abstractions and implementations used.

The rest of the article is organized as follows. Section 2 presents essential abstractions in logic languages. Sections 3, 4, and 5 describe abstractions, implementations, and applications centered around join, recursion, and constraint. Section 6 discusses additional language extensions, applications, and programming paradigms. Section 7 discusses related literature and future directions.

2 Logic language abstractions

Logic languages provide very high-level data and control abstractions, using mostly very simple language constructs. We describe these abstractions and their meanings intuitively.

2.1 Data abstractions

All data in logic languages are abstracted, essentially, as objects and relationships.

Objects. Objects are primitive values, such as numbers and strings, or structured values whose components are objects.

Examples of primitive values are integer number 3 and string 'Amy'. We enclose a string value in single quotes; if a string starts with a lower-case letter, such as 'amy', the quotes can be omitted, as has been conventional in logic languages.

Examples of structured values are `succ(3)`, `father(amy)`, and `cert('Amy', birth('2000-02-28', 'Rome'))`, denoting the successor integer of 3, the father of amy, and the certificate that 'Amy' was born on '2000-02-28' in 'Rome', respectively.

The names of structures, such as `succ`, `father`, `cert`, and `birth` above, are called function symbols. They correspond to object constructors in object-oriented languages.

Relationships. Relationships are predicates, or properties, that hold among objects. In particular, $p(o_1, \dots, o_k)$, i.e., predicate p over objects o_1, \dots, o_k being true, is equivalent to (o_1, \dots, o_k) in p , i.e., tuple (o_1, \dots, o_k) belonging to relation p —a table that holds the set of tuples of objects over which p is true.

Examples of relationships are `male(bob)`, `is_parent(bob, amy)`, and `issue(mario, 'Amy', birth('2000-02-28', 'Rome'))`, denoting that `bob` is male, `bob` is a parent of `amy`, and `mario` issued a certificate that `'Amy'` was born on `'2000-02-28'` in `'Rome'`, respectively.

Structured values can be easily captured using relationships, but not vice versa. For example, `f` being the structured value `father(c)` can be captured using relationship `is_father(f, c)`, but relationship `is_parent(p, c)` cannot simply be captured as `p` being the structured value `parent(c)` when `c` has two parents.

Such high-level data abstraction allows real-world objects or their lower-level representations, from bits and characters to lists to sets, to be captured easily without low-level implementation details. For example,

- bits and characters are special cases of integers and strings, respectively.
- lists are a special case of linearly nested structured values, and
- sets are a special case of relations consisting of tuples of one component.

Objects and relationships can be implemented using well-known data structures, including linked list, array, hash table, B-tree, and trie, usually taking $O(1)$ or $O(\log n)$ time to access an object, where n is the size of the data.

2.2 Control abstractions

Control in logic languages is abstracted at a high level, as logical inference or logic queries over asserted relationships among objects:

- asserted relationships can be connected by logical connectives: conjunction (read “and”), disjunction (read “or”), negation (read “not”), implication (read “then”), backward implication (read “if”), and equivalence (read “if and only if”);
- variables can be used in place of objects and be quantified over with universal quantifier (read “all”) and existential quantifier (read “some”); and
- one can either infer all relationships that hold or query about certain relationships, among all objects or among certain objects.

Rules and facts are the most commonly supported forms in existing logic languages:

Rules. A rule is of the following form, where *assertion₀* is called the conclusion, and other assertions are called the hypotheses. Each assertion is a predicate over certain objects, where variables may be used in place of objects. Intuitively, left arrow (<-) indicates backward implication, comma (,) denotes conjunction, and all variables in a rule are implicitly universally quantified, i.e., the rule holds for all values of the variables.

$$assertion_0 \leftarrow assertion_1, \dots, assertion_k.$$

For example, the second rule below says: *x* is a grandfather of *y* if *x* is the father of *z* and *z* is a parent of *y*, and this holds for all values of variables *x*, *y*, and *z*; the other rules can be read similarly. Following logic language conventions, names starting with an upper-case letter are variables.

```
is_parent(X,Y) <- is_father(X,Y).
is_grandfather(X,Y) <- is_father(X,Z), is_parent(Z,Y).
is_ancestor(X,Y) <- is_parent(X,Z), is_ancestor(Z,Y).
is_positive(succ(N)) <- is_positive(N).
```

The second rule is a join query—its two hypotheses have a shared variable, and it concludes a new predicate.

The third and fourth rules are recursive—the predicate in the conclusion depends on itself in a hypothesis, or in general possibly indirectly through another predicate.

Note that disjunction of a set of hypotheses can be expressed using a set of rules with the same conclusion.

Facts. A fact is a rule that has no hypotheses and is denoted simply as *assertion₀*. For example, *is_father(bob,amy).* says that *bob* is the father of *amy*, and *is_positive(1).* says that 1 is positive.

The meaning of a set of rules and facts is the least set of facts that contains all the given facts and all the facts that can be inferred, directly or indirectly, using the rules. This set can be computed by starting with the given facts and repeatedly applying the rules to conclude new facts—i.e., matching hypotheses of rules against facts, instantiating variables in rules with values in matched facts, and adding instantiated conclusions of rules as new facts. However,

- repeated application of rules might not terminate if function symbols are used in the rules, because facts about infinitely many new objects may be concluded, e.g., the fourth example rule above may infer *is_positive(succ(1)), is_positive(succ(succ(1))),* and so on.
- when only certain relationships about certain objects are queried, application of rules may stop as soon as the query can be answered, e.g., if only *is_positive(succ(1))* is queried, application of rules can stop after one use of the given rule and the given fact.

Rules that do not contain function symbols are called Datalog rules. For example, the first three example rules given earlier in this section are Datalog rules.

General logic forms have also been increasingly supported, typically by extending the rule form above:

Negation in the hypotheses. A hypothesis in a rule may be prefixed with `not`, denoting negation of the asserted relationship.

For example, the following rule says: for all values of `x` and `y`, `x` is the mother of `y` if `x` is a parent of `y` and `x` is not male.

```
is_mother(X,Y) <- is_parent(X,Y), not male(X).
```

Difficulties arise when negation is used with recursion. For example, what can be inferred from the following rule? Is `good(zak)` true or false?

```
good(zak) <- not good(zak).
```

More general forms. More general forms include disjunction and negation in the conclusion and, most generally, quantifiers `all` and `some` in any scope, not only the outermost scope. For example, the first rule below says: `x` is male or female if `x` is a person. The second rule says: `x` is not a winning position if, for all `y`, there is no move from `x` to `y` or else `y` is a winning position.

```
male(X) or female(X) <- person(X).
not win(X) <- all Y: not move(X,Y) or win(Y).
```

The meaning of recursive rules with negation is not universally agreed upon. The two dominant semantics are well-founded semantics (WFS) [VRS91, VG93] and stable model semantics (SMS) [GL88]. Both WFS and SMS use the closed-world assumption, i.e., they assume that what cannot be inferred to be true from the given facts and rules, is false.

- WFS gives a single 3-valued model, with the additional truth value `undefined` besides `true` and `false`.
- SMS gives zero or more 2-valued models, using only `true` and `false`.

Other formalisms and semantics include partial stable models, also called stationary models [Prz94]; first-order logic with inductive and fixed-point definitions, called FO(ID) and FO(FD) [DT08, HDCD10]; and the newly proposed founded semantics and constraint semantics [LS18]. The first two are both aimed at unifying WFS and SMS. The last unifies and cleanly relates WFS, SMS, and other major semantics by allowing the assumptions about the predicates and rules to be specified explicitly.

For practical applications, logic languages often also support predefined relationships among objects, including equality, inequality, and general comparisons. Cardinality and other aggregates over relationships are often also supported.

2.3 Combinations of control abstractions

There are many possible combinations of the language constructs. We focus on the following three combinations of constructs as essential control abstractions. We identify them by join, recursion, and constraint. They capture bounded, cyclic, and general computations, respectively.

- (1) **Join**—with join queries, no recursive rules, and restricted negation and other constructs; the restriction is that, for each rule, each variable in the conclusion must also appear in a hypothesis that is a predicate over arguments. Implementing this requires that common objects for the shared variables be found for the two hypotheses of a join query to be true at the same time; the number of objects considered are bounded, by the predicates in the hypotheses, following a bounded number of dependencies.
- (2) **Recursion**—with join queries, recursive rules, and restricted negation and other constructs; the restriction is as for join above plus that a predicate in the conclusion of a rule does not depend on the negation of the predicate itself in a hypothesis. Implementing this requires repeatedly applying the recursive rules following cyclic dependencies, potentially an unbounded number of times if new objects are in some conclusions.
- (3) **Constraint**—with join queries, recursive rules, and unrestricted negation and other constructs; unrestricted negation and other constructs can be viewed as constraints to be satisfied. Implementing this could require, in general, trying different combinations of variable values, as in general constraint solving.

Table 1 summarizes these three essential control abstractions and the corresponding kinds of computations and applications.

	Essential	Has join queries	Has rec. rules	Has neg. and others	Computations	Application kinds
(1)	Join	yes	no	restricted	bounded	database-style queries
(2)	Recursion	yes	yes	restricted	cyclic	inductive analysis
(3)	Constraint	yes	yes	unrestricted	general	combinatorial search

Table 1: Essential control abstractions of logic languages.

3 Join and database-style queries

Join queries are the most basic and most commonly used queries in relating different objects. They underlie essentially all nontrivial queries in database applications and many other applications.

3.1 Join queries

A join query is a conjunction of two hypotheses that have shared variables, concluding possible values of variables that satisfy both hypotheses. A conjunction of two hypotheses that have no shared variables, i.e., a Cartesian product, or a single hypothesis can be considered a trivial join query. A join query corresponds to a rule whose predicate in the conclusion is different from predicates in the hypothesis, so the rule is not recursive. A non-recursive rule with more than two hypotheses corresponds to multiple join queries, as a nesting or chain of join queries starting with joining any two hypotheses first.

For example, the first rule below, as seen before, is a join query. So is the second rule; it defines `sibling` over `x` and `y` if `x` and `y` have a same parent. The third rule defines a chain of red, green, and blue links from `x` to `y` through `u` and `v`; it can be viewed as two join queries—join any two hypotheses first, and then join the result with the third hypothesis.

```
is_grandfather(X,Y) <- is_father(X,Z), is_parent(Z,Y).
sibling(X,Y) <- is_parent(Z,X), is_parent(Z,Y).
chain(X,Y) <- link(X,U,red), link(U,V,green), link(V,Y,blue).
```

In general, the asserted predicates can be about relationships among any kinds of objects—whether people, things, events, or anything else, e.g., students, employees, patients, doctors, products, courses, hospitals, flights, interviews, and hangouts; and the join queries can be among any kinds of relationships—whether family, friend, owning, participating, thinking, or any other relation in the real world or conceptual world.

Join queries expressed using rules correspond to set queries. For example, in a language that supports set comprehensions with tuple patterns [RL07, LBSL16] the `is_grandfather` query corresponds to

```
is_grandfather = {(X,Y): (X,Z) in is_father, (Z,Y) in is_parent}
```

Without recursion, join queries can be easily supported together with the following extensions, with the restriction that, for each rule, each variable in the conclusion must also appear in a hypothesis that is a predicate over arguments, so the domain of the variable is bounded by the predicate; queries using these extensions can be arbitrarily nested:

- unrestricted negation, other connectives, and predefined relationships in additional conditions,
- aggregates, such as count and max, about the relationships, and
- general universal and existential quantifiers in any scope.

These subsume all constructs in the `select` statement for SQL queries. Essentially, join queries, with no recursion, relate objects in different relationships within a bounded number of steps.

3.2 Implementation of join queries

A join query can be implemented straightforwardly using nested for-loops and if-statements, where shared variables in different hypotheses correspond to equality tests between the corresponding variables. For example, the `is_grandfather` query earlier in this section can be implemented as

```
is_grandfather = {}
for (X,Z1) in is_father:           -- time factor: number of is_father pairs
    for (Z2,Y) in is_parent:       -- time factor: number of is_parent pairs
        if Z1 == Z2:
            is_grandfather.add(X,Y)
```

In a language that supports set comprehensions, such as Python, the above implementation can be expressed as

```
is_grandfather = {(X,Y) for (X,Z1) in is_father for (Z2,Y) in is_parent if Z1 == Z2}
```

For efficient implementations, several key implementation and optimization techniques are needed, described below; additional optimizations are also needed, e.g., for handling streaming data or distributed data.

Indexing. This creates an index for fast lookup based on values of the indexed arguments of a relation; the index is on the shared arguments of the two hypotheses. For example, for any fact `is_father(X,Z)`, to find the matching `is_parent(Z,Y)`, an index called, say, `children{Z}`—mapping the value of `Z`, the first argument of `is_parent`, to the set of corresponding values of second argument of `is_parent`—significantly speeds up the lookup, improving the time factor for the inner loop to the number of children of `Z`:

```
is_grandfather = {(X,Y) for (X,Z) in is_father for Y in children{Z}}
```

Join ordering. This optimizes the order of joins when there are multiple joins, e.g., in a rule with more than two hypotheses. For example, for the rule for `chain`, starting by joining the first and third hypotheses is never more efficient than starting by joining either of these hypotheses with the second hypothesis, because the former yields all pairs of red and blue links, even if there are no green links in the middle.

Tabling. This stores the result of common sub-joins so they are not repeatedly computed. Common sub-joins may arise when there are nested or chained join queries. For example, for the rule for `chain` earlier in this section, consider joining the first two hypotheses first: if there are many red and green link pairs from a value of `x` to a value of `v`, then storing the result of this sub-join avoids recomputing it when joining with blue links to find each target `y`.

Demand-driven computation. This computes only those parts of relationships that affect a particular query. For example, a query may only check whether `is_father(dan, bob)` holds, or find all values of `x` for `is_father(dan, x)`, or find all `is_father` pairs, as opposed to finding all relationships that can be inferred.

Basic ideas for implementing the extensions negation, aggregates, etc. are as follows, where nested queries using these extensions are computed following their order of dependencies:

- negation, etc. in additional conditions: test them after the variables in them become bound by the joins.
- aggregates: apply the aggregate operation while collecting the query result of its argument.
- quantifiers: transform them into for-loops, or into aggregates, e.g., an existential quantification is equivalent to a count being positive.

Efficient implementation techniques for join queries and extensions have been studied in a large literature, e.g., [Ioa96]. Some methods also provide precise complexity guarantees, e.g., [Wil02, LBSL16].

3.3 Applications of join queries

Join queries are fundamental in querying complex relationships among objects. They are the core of database applications [KBL06], from enterprise management to ontology management, from accounting systems to airline reservation systems, and from electronic health records management to social media management. Database and logic programming are so closely related that one of the most important computer science bibliographies is called DBLP, and it was named after Database and Logic Programming [Ley02]. Join queries also underlie applications that do not fit in traditional database applications, such as complex access control policy frameworks [ANS04].

We describe three example applications below, in the domains of ontology management, enterprise management, and security policy frameworks. They all heavily rely on the use of join queries and optimizations, especially indexing. We give specific examples of facts, rules, and indexing for the first application.

Ontology management—Coherent definition framework (CDF). CDF is a system for ontology management that has been used in numerous commercial projects [GAS10], for organizing information about, e.g., aircraft parts, medical supplies, commercial processes, and materials. It was originally developed by XSB, Inc. Significant portions have been released in the XSB packages [SW⁺14].

The data in CDF are classes and objects. For example, XSB, Inc. has a part taxonomy, combining UNSPSC (United Nations Standard Products and Services Code) and Federal INC (Item Name Code) taxonomies, with a total of over 87,000 classes of parts. The main relationships are variants of `isa`, `hasAttr`, and `allAttr`. Joins are used extensively to answer queries about closely related classes, objects, and attributes. Indexing and tabling are heavily used for efficiency. Appropriate join order and demand-driven computation are also important.

An example fact is as follows, indicating that specification 'A-A-1035' in ontology `specs` has attribute 'MATERIAL' whose value is 'ALUMINUM ALLOY UNS A91035' in `material_taxonomy`. Terms `cid(Identifier, Namespace)` represent primitive classes in CDF.

```
hasAttr_MATERIAL(cid('A-A-1035',specs),
                 cid('ALUMINUM ALLOY UNS A91035',material_taxonomy)).
```

An example rule is as follows, meaning that a part `PartNode` has attribute 'PART-PROCESS-MATERIAL' whose value is process-material pair (`Process`,`Material`) in 'ODE Ontology' if `PartNode` has attribute 'PROCESS' whose value is `Process`, and `Process` has attribute 'PROCESS-MATERIAL' whose value is `Material`.

```
hasAttr_PART-PROCESS-MATERIAL(PartNode,
                               cid('process-material'(Process,Material),'ODE Ontology')) <-
  hasAttr_PROCESS(PartNode, Process),
  hasAttr_PROCESS-MATERIAL(Process, Material).
```

An example of indexing is for `hasAttr_ATTR`, for any `ATTR`, shown below, in XSB notation, meaning: use as index all symbols of the first argument if it is bound, or else do so for the second argument.

```
[*(1), *(2)]
```

XSB, Inc. has five major ontologies represented in CDF, for parts, materials, etc., with a total of over one million facts and five meta rules. The rules are represented using a Description Logic form—an ontology representation language. The example rule above is an instance of such a rule when interpreted. The indexing used supports different appropriate indices for different join queries.

CDF is used in XSB, Inc.'s ontology-directed classifier (ODC) and extractor (ODE) [SW12]. ODC uses a modified Bayes classifier to classify item descriptions. For example, it is used quarterly by the U.S. Department of Defense to classify over 80 million part descriptions. ODE extracts attribute-value pairs from classified descriptions to build structured knowledge about items. ODC uses aggregates extensively, and ODE uses string pattern rules.

Enterprise management—Business intelligence (BI). BI is a central component of enterprise software. It tracks the performance of an enterprise over time by storing and analyzing historical information recorded through online transaction processing (OLTP), and is then used to help plan future actions of the enterprise. LogicBlox simplifies the hairball of enterprise software technologies by using a Datalog-based language [GAK12, AtCG⁺15].

All data are captured as logic relations. This includes not only data as in conventional databases, e.g., sale items, price, and so on for a retail application, but also data not in conventional databases, e.g., sale forms, display texts, and submit buttons in a user interface. Joins are used for easily querying interrelated data, as well as for generating user interfaces. Many extensions such as aggregates are also used. For efficiency, exploiting the rich literature of automatic optimizations, especially join processing strategies and incremental maintenance, is of paramount importance.

Using the same Datalog-based language, LogicBlox supports not only BI but also OLTP and prescriptive and predictive analytics. “Today, the LogicBlox platform has matured to the point that it is being used daily in dozens of mission-critical applications in some of the largest enterprises in the world, whose aggregate revenues exceed \$300B” [AtCG⁺15].

Security policy frameworks—Core role-based access control (RBAC). RBAC is a framework for controlling user access to resources based on roles. It became an ANSI standard [ANS04] building on much research during the preceding decade and earlier, e.g., [LHM84, FK92, GB98, FSG⁺01].

Core RBAC defines users, roles, objects, operations, permissions, sessions and a number of relations among these sets; the rest of RBAC adds a hierarchical relation over roles, in hierarchical RBAC, and restricts the number of roles of a user and of a session, in constrained RBAC. Join queries are used for all main system functions, especially the `CheckAccess` function, review functions, and advanced review functions on the sets and relations. They are easily expressed using logic rules [BLV04, BF06].

Efficient implementations rely on all main optimizations discussed, especially auxiliary maps for indexing and tabling [LWG⁺06]. Although the queries are like relational database queries, existing database implementations would be too slow for functions like `CheckAccess`. Unexpectedly, uniform use of relations and join queries also led to a simplified specification, with unnecessary mappings removed, undesired omissions fixed, and constrained RBAC drastically simplified [LS07].

4 Recursion and inductive analysis

Recursive rules are most basic and essential in relating objects that are an unknown number of relationships apart. They are especially important for problems that may require

performing the inference or queries a non-predetermined number of steps, depending on the data.

4.1 Recursive rules and queries

Given a set of rules, a predicate p depends on a predicate q if p is in the conclusion of a rule, and either q is in a hypothesis of the rule or some predicate r is in a hypothesis of the rule and r depends on q . A given set of rules is recursive if a predicate p in the conclusion of a rule depends on p itself.

For example, the second rule below, as seen in Section 2.2, is recursive; the first rule is not recursive; the set of these two rules is recursive, where the first rule is the base case, and the second rule is the recursive case.

```
is_ancestor(X,Y) <- is_parent(X,Y).
is_ancestor(X,Y) <- is_parent(X,Z), is_ancestor(Z,Y).
```

In general, recursively asserted relationships can be between objects of any kind, e.g., relatives and friends that are an unknown number of connections apart in social networks, direct and indirect prerequisites of courses in universities, routing paths in computer networks, nesting of parts in products, supply chains in supply and demand networks, transitive role hierarchy relation in RBAC, and repeated delegations in trust management systems.

Recursive queries with restricted negation correspond to least fixed-point computations. For example, in a language that supports least fixed points, the `is_ancestor` query corresponds to the minimum `is_ancestor` set below, where, for any sets S and T , $S \text{ subset } T$ holds iff every element of S is an element of T :

```
min is_ancestor: is_parent subset is_ancestor,
                 {(X,Y): (X,Z) in is_parent, (Z,Y) in is_ancestor} subset is_ancestor
```

With cyclic predicate dependencies, recursion allows the following restricted extensions to be supported while still providing a unique semantics; there is also the restriction that, for each rule, each variable in the conclusion must also appear in a hypothesis that is a predicate over arguments, as in extensions to join queries:

- stratified negation, where negation and recursion are separable, i.e., there is no predicate that depends on the negation of itself, and
- other connectives and predefined relationships in additional conditions, aggregates, and general quantifiers, as in extensions for join queries, when they do not affect the stratification.

Essentially, recursive rules capture an unbounded number of joins, and allow inference and queries by repeatedly applying the rules.

4.2 Implementation of recursive rules and queries

Inference and queries using recursive rules can be implemented using while-loops; for-loops with predetermined number of iterations do not suffice, because the number of iterations depends on the rules and facts. Each iteration applies the rules in one step, so to speak, until no more relevant facts can be concluded. For example, the `is_ancestor` query earlier in this section can be implemented as

```
is_ancestor = is_parent
while exists (X,Y): (X,Z) in is_parent, (Z,Y) in is_ancestor, (X,Y) not in is_ancestor:
    is_ancestor.add((X,Y))
```

Each iteration computes the existential quantification in the condition of the while-loop, and picks any witness (X,Y) to add to the result set. It can be extremely inefficient to recompute the condition in each iteration after a new pair is added.

For efficient implementations, all techniques for joins are needed but are also more critical and more complex. In particular, to ensure termination,

- tabling is critical if relationships form cycles, and
- demand-driven computation is critical if new objects are created in the cycles.

For the `is_ancestor` query, each iteration computes the following set, which is a join, plus the last test to ensure that only a new fact is added:

```
{(X,Y): (X,Z) in is_parent, (Z,Y) in is_ancestor, (X,Y) not in is_ancestor}
```

Two general principles underlying the optimizations for efficient implementations are:

1. incremental computation for expensive relational join operations, with respect to facts that are added in each iteration.
2. data structure design for the relations, for efficient retrievals and tests of relevant facts.

For the restricted extensions, iterative computation follows the order of dependencies determined by stratification; additional aggregates, etc. that do not affect the stratification can be handled as described in Section 3.2 for computing the join in each iteration.

Efficient implementation techniques for recursive queries and extensions have been studied extensively, e.g., [AHV95]. Some methods also provide precise complexity guarantees, e.g., [McA99, GM01, LS09, TL11].

4.3 Applications of recursive rules and queries

Recursive rules and queries can capture any complex reachability problem in recursive structures, graphs, and hyper-graphs. Examples are social network analysis based on all kinds of social graphs; program analysis over many kinds of flow and dependence graphs about program control and data values; model checking over labeled transition systems and state machines; routing in electronic data networks, telephone networks, or transportation networks; and security policy analysis and enforcement over trust or delegation relationships.

We describe three example applications below, in the domains of text and natural language processing, program analysis, and distributed security policy frameworks. They all critically depend on the use of recursive rules and efficient implementation techniques, especially tabling and indexing.

Text processing—Super-tokenizer. Super-tokenizer is an infrastructure tool for text processing that has been used by XSB, Inc.’s ontology-directed classifier (ODC) and extractor (ODE) for complex commercial applications [SW12]. It was also developed originally at XSB, Inc.

Super-tokenizer supports the declaration of complex rewriting rules for token lists. For example, over 65,000 of these rules implement abbreviations and token corrections in ODC and complex pattern-matching rules in ODE for classification and extraction based on combined UNSPSC and Federal INC taxonomies at XSB, Inc. Recursion is used extensively in the super-tokenizer, for text parsing and processing. The implementation uses tabled grammars and trie-based indexing in fundamental ways.

Super-tokenizer is just one particular application that relies on recursive rules for text processing and, more generally, language processing. Indeed, the original application of Prolog, the first and main logic programming language, was natural language processing (NLP) [PS02], and a more recent application in NLP helped the IBM Watson question answering system win the Jeopardy Man vs. Machine Challenge by defeating two former grand champions in 2011 [LF11, LPM⁺12].

Program analysis—Pointer analysis. Pointer analysis statically determines the set of objects that a pointer variable or expression in a program can refer to. It is a fundamental program analysis with wide applications and has been studied extensively, e.g., [Hin01, SCD⁺13]. The studies especially include significantly simplified specifications using Datalog in more recent years, e.g., [SB15], and powerful systems such as bddb [WACL05] and Doop [BS09b], the latter built using LogicBlox [GAK12, AtCG⁺15].

Different kinds of program constructs and analysis results relevant to pointers are relations. Datalog rules capture the analysis directly as recursively defined relations. For example, the well-known Andersen’s pointer analysis for C programs defines a points-to relation based on four kinds of assignment statements [And94], leading directly to four Datalog rules [SR05]. Efficient implementation critically depends on tabling, indexing,

and demand-driven computation [SR05, TL11]. Such techniques were in fact followed by hand to arrive at the first ultra fast analysis [HT01b, HT01a].

Indeed, efficient implementations can be generated from Datalog rules giving much better, more precise complexity guarantees [LS09, TL11] than the worst-case complexities, e.g., the well-known cubic time for Andersen’s analysis. Such efficient implementation with complexity guarantees can be obtained for program analysis in general [McA99]. Commercial tools for general program analysis based on Datalog have also been built, e.g., by Semmler based on CodeQuest [HVM06].

Security policy frameworks—Trust management (TM). TM is a unified approach to specifying and enforcing security policies in distributed systems [BFL96, GS00, RK05]. It has become increasingly important as systems become increasingly interconnected, and logic-based languages have been used increasingly for expressing TM policies [Bon10], e.g., SD3 [Jim01], RT [LMW02], Binder [DeT02], Cassandra [BS04], and many extensions, e.g., [BRS12, SBK13].

Certification, delegation, authorization, etc. among users, roles, permissions, etc. are relations. Policy rules correspond directly to logic rules. The relations can be transitively defined, yielding recursive rules. For example, one of the earliest TM frameworks, SPKI/SDSI [EFL⁺99], for which various sophisticated methods have been studied, corresponds directly to a few recursive rules [HTL07], and efficient implementations with necessary indexing and tabling were generated automatically.

TM studies have used many variants of Datalog with restricted constraints [LM03], not unrestricted negation. A unified framework with efficient implementations is still lacking. For example, based on the requirements of the U.K. National Health Service, a formal electronic health records (EHR) policy was written, as 375 rules in Cassandra [Bec05b], heavily recursive. As the largest case study in the TM literature, its implementation was inefficient and incomplete—techniques like indexing were deemed needed but missing [Bec05a].

5 Constraint and combinatorial search

Constraints are the most general form of logic specifications, which easily captures the most challenging problem-solving activities such as planning and resource allocation.

5.1 Constraint satisfaction

A constraint is, in general, a relationship among objects but especially refers to cases when it can be satisfied with different choices of objects and the right choice is not obvious.

For example, the rule below says that x is a winning position if there is a move from x to y and y is not a winning position. It states a relationship among objects, but its meaning

is not obvious, because the concluding predicates are recursively defined using a negation of the predicate itself.

```
win(X) <- move(X,Y), not win(Y).
```

In general, constraints can capture any real-world or conceptual-world problems, e.g., rules for moves in any game—whether recreational, educational, or otherwise; actions with conditions and effects for any planning activities; participants and resource constraints in scheduling—whether for university courses or manufacturer goods production or hospital surgeries; real-world constraints in engineering design; as well as knowledge and rules for puzzles and brain teasers.

Given constraints may have implications that are not completely explicit. For example, the `win` rule implies not just the first constraint below, but also the second, by negating the conclusion and hypotheses in the given rule, following the closed-world assumption; the second constraint makes the constraint about `not win` explicit:

```
win(X) if some Y: move(X,Y) and not win(Y)
not win(X) if all Y: not move(X,Y) or win(Y)
```

Indeed, with general constraints, objects can be related in all ways using all constructs together with join and recursion: unrestricted negation, other connectives, predefined relationships, aggregates, and general quantifiers in any scope.

However, due to negation in dependency cycles, the meaning of the rules and constraints is not universally agreed on anymore.

- Well-founded semantics (WFS) gives a single, 3-valued model, where relationships that are true or false are intended to be supported from given facts, i.e., well-founded, and the remaining ones are `undefined`.
- Stable model semantics (SMS) gives zero or more 2-valued models, where each model stays the same, i.e., is stable, when it is used to instantiate all the rules; in other words, applying the rules to each model yields the same model.

For example, for the `win` example,

- if there is only one move, `move(a,b)`, not forming a cycle, then
WFS and SMS both give that `win(b)` is false and `win(a)` is true;
- if there is only one move, `move(a,a)`, forming a self cycle, then
WFS gives that `win(a)` is undefined, and
SMS gives that there is no model;

- if there are only two moves, `move(a,b)` and `move(b,a)`, forming a two-move cycle, then WFS gives that `win(a)` and `win(b)` are both undefined, and SMS gives two models: one with `win(a)` true and `win(b)` false, and one with the opposite results.

Despite the differences, WFS and SMS can be computed using some shared techniques.

5.2 Implementation of constraint satisfaction

Constraint solving could in general use straightforward generate-and-test—generate each possible combination of objects for solutions and test whether they satisfy the constraints—but backtracking is generally used, as it is much more efficient.

Backtracking. Backtracking incrementally builds variable assignments for the solutions, and abandons each partial assignment as soon as it determines that the partial assignment cannot be completed to a satisfying solution, going back to try a different value for the last variable assigned; this avoids trying all possible ways of completing those partial assignments or naively enumerating all complete assignments.

For example, the `win(X)` query can basically try a move at each next choice of moves and backtrack to try a different move as soon as the current move fails. Expressed using recursive functions, this corresponds basically to the following:

```
def win(X): return (some Y: move(X,Y) and not_win(Y))
def not_win(X): return (all Y: not move(X,Y) or win(Y))
```

This backtracking answers the query correctly when the moves do not form a cycle. However, it might not terminate when the moves form a cycle, and the implementation depends on the semantics used. Both WFS and SMS can be computed by using and extending the basic backtracking:

- WFS computation could track cycles, where executing a call requires recursively making the same call, and infer undefined for those queries that have no execution paths to infer the query result to be true or false.
- SMS computation could generate possible partial or complete variable assignments, called grounding, and check them, possibly with the help of an external solver like Boolean satisfiability (SAT) solvers or satisfiability modulo theories (SMT) solvers.

For efficient implementations, techniques for join and recursion are critical as before, especially tabling to avoid repeated states in the search space. Additionally, good heuristics for pruning the search space can make drastic performance difference in computing SMS, e.g., as implemented in answer set programming (ASP) solvers.

Backjumping. One particular optimization of backtracking in SMS computation is backjumping. Backtracking always goes back one level in the search tree when all values for a variable have been tested. Backjumping may go back more levels, by realizing that a prefix of the partial assignment can lead to all values for the current variable to fail. This helps prune the search space.

For extensions that include additional constraints, such as integer constraints, as well as aggregates and quantifiers, an efficient solver such as one that supports mixed integer programming (MIP) can be used.

Efficient implementation techniques for constraint solving have been studied extensively, e.g., for ASP solvers [LPF⁺06, GKKS12].

5.3 Applications of constraint satisfaction

The generality and power of constraints allow them to be used for all applications described previously, but constraints are particularly important for applications beyond those and that require combinatorial search. Common kinds of search problems include planning and scheduling, resource allocation, games and puzzles, and well-known NP-complete problems such as graph coloring, k-clique, set cover, Hamiltonian cycle, and SAT.

We describe three example applications, in the domains of decision making, resource allocation, and games and puzzles. They all require substantial use of general constraints and efficient constraint solvers exploiting backtracking, backjumping, and other optimizations.

Enterprise decision making—Prescriptive analysis. Prescriptive analysis suggests decision options that lead to optimized future actions. It is an advanced component of enterprise software. For example, for planning purposes, LogicBlox supports prescriptive analysis using the same Datalog-based language as for BI and OLTP [GAK12, AtCG⁺15].

The data are objects and relations, same as used for BI, but may include, in particular, costs and other objective measures. Constraints capture restrictions among the objects and relations. When all data values are provided, constraints can simply be checked. When some data values are not provided, different choices for those values can be explored, and values that lead to certain maximum or minimum objective measures may be prescribed for deciding future actions. Efficient implementations can utilize the best constraint solvers based on the kinds of constraints used.

LogicBlox’s integrated solution to decision making based on BI and OLTP has led to significant success. For example, for a Fortune 50 retailer with over \$70 billion in revenue and with products available through over 2,000 stores and digital channels, the solution processes 3 terabytes of data on daily, weekly, and monthly cycles, deciding exactly what products to sell in what stores in what time frames; this reduces a multi-

year cycle of a challenging task for a large team of merchants and planners to an automatic process and significantly increases profit margins [Log15a].

Resource allocation—Workforce management (WFM) in Port of Gioia Tauro. WFM handles activities needed to maintain a productive workforce. The WFM system for automobile logistics in the Port of Gioia Tauro, the largest transshipment terminal in the Mediterranean, allocates available personnel of the seaport such that cargo ships mooring in the port are properly handled [RGA⁺12, LR15]. It was developed using the DLV system [LPF⁺06].

The data include employees of different skills, cargo ships of different sizes and loads, teams and roles to be allocated, and many other objects to be constrained, e.g., workload of employees, heaviness of roles, and contract rules. Constraints include matching of available and required skills, roles, hours, etc., fair distribution of workload, turnover of heavy or dangerous roles, and so on. The constraints are expressed using rules with disjunction in the conclusion, general negation, and aggregates. The DLV system uses backtracking and a suite of efficient implementation techniques.

This WFM system was developed by Exeura s.r.l. and has been adopted by the company ICO BLG operating automobile logistics in the Port of Gioia Tauro [LR15], handling every day several ships of different sizes that moor in the port [RGA⁺12].

Games and puzzles—N-queens. We use a small example in a large class of problems. The n-queens puzzle is the problem of placing n queens on a chessboard of n -by- n squares so that no two queens threaten each other, i.e., no two queens share the same row, column, or diagonal. The problem is old, well-studied, and can be computationally quite expensive [BS09a].

The allowed placements of queens can be specified as logic rules with constraints. Naively enumerating all possible combination of positions and checking the constraints is prohibitively expensive. More efficient solutions use backtracking, and furthermore backjumping, to avoid impossible placement of each next queen as soon as possible. Stronger forms of constraints may also be specified to help prune the search space further [GKKS12]. For example, backtracking can solve for one or two scores of queens in an hour, but backjumping and additional constraints help an ASP system like Clingo solve for 5000 queens in 3758.320 seconds of CPU time [Sch14].

Many other games and puzzles can be specified and solved in a similar fashion. Examples are all kinds of crossword puzzles, Sudoku, Knight’s tour, nonograms, magic squares, dominos, coin puzzles, graph coloring, palindromes, among many others, e.g., [DNST05, Edm15, Het15, Mal15, Kje15].

6 Further extensions, applications, and discussion

We discuss additional language extensions and applications, summarize applications based on the key abstractions used, touching on example logic programming systems, and finally put the abstractions into the perspective of programming paradigms.

6.1 Extensions

Many additional extensions to logic languages have been studied. Most of them can be viewed as abstractions that capture common patterns in classes of applications, to allow applications to be expressed more easily. Important extensions include:

- regular-expression paths, a higher-level abstraction for commonly-used linear recursion;
- updates, for real-world applications that must handle changes;
- time, for expressing changes over time, as an alternative to supporting updates directly;
- probability, to capture uncertainty in many challenging applications; and
- higher-order logic, to support applications that require meta-level reasoning.

We discuss two of the most important extensions below:

Regular-expression paths. A regular-expression path relates two objects using regular expressions and extensions. It allows repeated joins of a binary relation to be expressed more easily and clearly than using recursion; such joins capture reachability and are commonly used. For example, `is_ancestor(X,Y)`, defined in Section 4 using two rules including a recursive rule, can now be defined simply as below; it indicates that there are one or more `is_parent` relationships in a path from `x` to `y`:

```
is_ancestor(X,Y) <- is_parent+(X,Y).
```

This is also higher-level than using recursion, because the recursive rule has to pick one of three possible forms below: with `is_parent` on the left, as seen before; with `is_parent` on the right; and with both conjuncts using `is_ancestor`.

```
is_ancestor(X,Y) <- is_parent(X,Z), is_ancestor(Z,Y).  
is_ancestor(X,Y) <- is_ancestor(X,Z), is_parent(Z,Y).  
is_ancestor(X,Y) <- is_ancestor(X,Z), is_ancestor(Z,Y).
```

Depending on the data, the performance of these forms can be asymptotically different in most implementations.

Regular-expression paths have many important applications including all those in Section 4, especially graph queries, with also parametric extensions for more general relations, not just binary relations [dMLW03, LRY⁺04, LS06, TGL10].

Updates. An update, or action, can be expressed as a predicate that captures the update, e.g., by relating the values before and after the update and the change in value. The effect of the update could be taken immediately after the predicate is evaluated, similar to updates in common imperative languages, but this leads to lower-level control flows that are harder to reason about. Instead, it is better for the update to take effect as part of a transaction of multiple updates that together satisfy high-level logic constraints. For example, with this approach, the following rule means that `adopted_by_from` holds if the updates `add_child` and `del_child` and the check `adoption_check` happen as a transaction.

```
adopted_by_from(C,X,Y) <- add_parent(X,C), del_parent(Y,C), adoption_check(C,X,Y).
```

It ensures at a high-level that certain bad things won't happen, e.g., no child would end with one fewer parent or one more parent than expected. Transaction logic is an extension of logic rules for reasoning about and executing transactional state changes [BK94]. LPS, a Logic-based approach to Production Systems, captures state changes by associating timestamps with facts and events, and this is shown to correspond to updating facts directly [KS15].

Logic languages with updates have important applications in enterprise software [GAK12, AtCG⁺15]. Transaction logic can also help in planning [BKB14].

Additional implementation support can help enhance applications and enable additional applications. A particular helpful feature is to record justification or provenance information during program execution [RRR00, DAA13], providing explanations for how a result was obtained. The recorded information can be queried to improve understanding and help debugging.

6.2 Additional applications

Many additional applications have been developed using logic programming, especially including challenging applications that need recursion and those that furthermore need constraint.

Table 2 lists example application areas with example application problems organized based on the main abstractions used. Note that application problems can often be reduced to each other, and many other problems can be reduced to the problems in the table. For example, model checking a property of a system [CGP99, CGL94] can be reduced to planning, where the goal state is a state violating the property specified, so a plan found by a planner corresponds to an error trace found by a model checker [EMP14, CEFP14]. Administrative policy analysis also has correspondences to planning, by finding a sequence of actions to achieve the effect of a security breach [SYGR11].

Table 2 is only a small sample of the application areas, with example application problems or kinds of application problems in those areas. Many more applications have been developed,

Area	Using join	Using recursion	Using constraint
Data management	business intelligence,* many database join queries	route queries, many database recursive queries	data cleaning, data repair
Knowledge management	ontology management*	ontology analysis	reasoning with knowledge
Decision support		supply-chain management, market analysis	prescriptive analysis,* planning, scheduling, resource allocation*
Linguistics		text processing,* context-free parsing, semantic analysis	context-sensitive analysis, deep semantics analysis
Program analysis	type checking, many local analyses	pointer analysis,* type analysis, many dependency analyses	type inference, many constraint-based analyses
Security	role-based access control*	trust management,* hierarchical role-based access control	administrative policy analysis, cryptanalysis
Games and puzzles		Hanoi tower, many recursion problems	n-queens,* Sudoku, many constraint puzzles
Teaching	course management	course analysis	question analysis, problem diagnosis, test generation

Table 2: Example application areas with example application problems organized based on the main abstractions used. Applications discussed in some detail in this article are marked with an asterisk.

in many more areas, using systems that support variants of the abstractions with different implementations. Some examples are:

- XSB has also been used to develop applications for immunization survey [BKGD⁺12], standardizing data, spend analysis, etc. [XSB15], and it is discussed in many publications¹.
- LogicBlox has also been used to create solutions for predicting consumer demand, optimizing supply chain, etc. [Log15b] and more [GAK12].
- ASP systems have been used in bioinformatics, hardware design, music composition, robot control, tourism, and many other application areas [Gro05, Sch11], including part

¹A Google Scholar search with +XSB +”logic programming” returns over 2300 results, July 2, 2017.

of a decision support system for the Space Shuttle flight controller [NBG⁺01, BG05].

- Logic systems have been developed for additional applications, e.g., PRISM [SK97] and ProbLog [DRKT07] for probabilistic models; XMC [RRS⁺00] and ProB [LB08] for verification; and NDlog [LCG⁺09], Meld [ARLG⁺09], Overlog [ACC⁺10], and Bloom [Ber13] for network and distributed algorithms.

Languages and systems with more powerful features such as constraints for general applications are often also used in less challenging application areas such as those that need only join queries. For example, DLV has also been used in ontology management [RGS⁺09].

6.3 Additional discussion on abstractions

We give an overview of the main abstractions in the larger picture of programming paradigms, to help put the kinds of applications supported into broader perspective.

The three main abstractions—join, recursion, and constraint—correspond generally to more declarative programming paradigms. Each is best known in its corresponding main programming community:

- Join in database programming. Database systems have join at the core but support restricted recursion and constraints in practice.
- Recursion in functional programming. Functional languages have recursion at the core but do not support high-level join or constraints.
- Constraint in logic programming. Logic engines support both join and recursion at the core, and have increasingly supported constraints at the core as well.

The additional extensions help further raise the level of abstraction and broaden the programming paradigms supported:

- Regular-expression paths raise the level of abstraction over lower-level linear recursion.
- Updates, or actions, are the core of imperative programming; they help capture real-world operations even when not used in low-level algorithmic steps.
- Time, probability, higher-order logic, and many other features correspond to additional arguments, attributes, or abstractions about objects and relationships.

One main paradigm not yet discussed is object-oriented programming. Orthogonal to data and control abstractions, objects in common object-oriented languages provide a kind of module abstraction, encapsulating both data structures and control structures in objects and classes. Similar abstractions have indeed been added to logic languages as well. For

example, F-logic extends traditional logic programming with objects [KLW95] and is supported in Flora-2 [KYWZ14]; it was also the basis of a highly scalable commercial system, Ontobroker [Sem12], and a recent industry suite, Ergo [GBF⁺15]. For another example, ASP has been extended with object constructs in OntoDLV [RGS⁺09].

Finally, building practical applications requires powerful libraries and interfaces for many standard functionalities. Many logic programming systems provide various such libraries. For example, SWI-Prolog has libraries for constraint logic programming, multithreading, interface to databases, GUI, a web server, etc, as well as development tools and extensive documentation.

7 Related literature and future work

There are many overview books and articles about logic programming in general and applications of logic programming in particular. This article differs from prior works by studying the key abstractions and their implementations as the driving force underlying vastly different application problems and application areas.

Kowalski [Kow14] provides an extensive overview of the development of logic programming. It describes the historical root of logic programming, starting from resolution theorem proving; the procedural interpretation and semantics of rules with no negated hypotheses, called Horn clause programs; negation as failure, including completion semantics, stratification, well-founded semantics, stable model semantics, and ASP; as well as logic programming involving abduction, constraints, and argumentation. It focuses on three important issues: logic programming as theorem proving vs. model generation, with declarative vs. procedural semantics, and using top-down vs. bottom-up computation. Our description of abstractions and implementations aims to separate declarative semantics from procedural implementations.

Other overviews and surveys about logic programming in general include some that cover a collection of topics together and some that survey different topics separately. Example collections discuss the first 25 years of logic programming from 1974 [AWT99] and the first 25 years of the Italian Association of Logic Programming from 1985 [DP10]. Example topics surveyed separately include logic programming semantics [Fit02], complexity and expressive power [DEGV01], constraints [JM94], ASP and DLV [GLR13], deductive databases [CGT90, AHV95, RU95, MSZ14], and many more. Our description of abstractions and implementations is only a highly distilled overview of the core topics.

Overviews and surveys about logic programming applications in particular are spread across many forums. Example survey articles include an early article on Prolog applications [Rot93], DLV applications [GILR09, GLMR11, LR15], applications in Italy [DPT10], emerging applications [HGL11], and a dedicated workshop AppLP—Applications of Logic Programming [WL17]. For example, the early article [Rot93] describes six striking practical applications of Prolog that replaced and drastically improved over systems written previously

using Fortran, C++, and Lisp. Example collections of applications on the Web include one at TU Wien [Gro05], one by Schaub [Sch11], and some of the problems in various competitions, e.g., as described by Gebser et al. [GMR17]. We try to view the applications by the abstractions and implementations used, so as to not be distracted by specific details of very different applications.

There are also many articles on specific applications or specific classes of applications. Examples of the former include team building [RGA⁺12], program pointer analysis [SB15], and others discussed in this article. Examples of the latter include applications in software engineering [CS95], DLV applications in knowledge management [GILR09], and IDP applications in data mining and machine learning [BBB⁺14]. We used a number of such specific applications as examples and described some of them in slightly more detail to illustrate the common technical core in addition to the applications per se.

Directions for future work. There are several main areas for future study: (1) more high-level abstractions that are completely declarative, (2) more efficient implementations with complexity guarantees, and (3) more unified and standardized languages and frameworks with rich libraries. These will help many more applications to be created in increasingly complex problem domains.

Acknowledgment

I would like to thank David S. Warren for his encouragement over the years at Stony Brook, and his patient and stimulating explanations about logic programming in general and XSB implementation in particular. I am grateful to Molham Aref, Francesco Ricca, and David Warren for helpful suggestions and additional information about applications using LogicBlox, DLV, and XSB, respectively. I thank Molham Aref and others at LogicBlox, Jon Brandvein, Christopher Kane, Michael Kifer, Bob Kowalski, Bo Lin, Francesco Ricca, Scott Stoller, Tuncay Tekle, David Warren, Neng-Fa Zhou, and anonymous reviewers for helpful comments on drafts of this article.

References

- [ACC⁺10] P. Alvaro, T. Condie, N. Conway, J.M. Hellerstein, and R. Sears. I do declare: Consensus in a logic language. *ACM SIGOPS Operating Systems Review*, 43(4):25–30, 2010.
- [AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases: The Logical Level*. Addison-Wesley, 1995.
- [And94] Lars Ole Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, University of Copenhagen, 1994.
- [ANS04] ANSI INCITS. Role-Based Access Control. ANSI INCITS 359-2004, American National Standards Institute, International Committee for Information Technology Standards, Feb. 2004.

- [ARLG⁺09] Michael P. Ashley-Rollman, Peter Lee, Seth Copen Goldstein, Padmanabhan Pillai, and Jason D. Campbell. A language for large ensembles of independently executing nodes. In *Proceedings of the 25th International Conference on Logic Programming*, pages 265–280. Springer, 2009.
- [AtCG⁺15] Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and Geoffrey Washburn. Design and implementation of the LogicBlox system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1371–1382, 2015.
- [AWT99] Krzysztof R. Apt, David S. Warren, and Mirek Truszczyński, editors. *The Logic Programming Paradigm: A 25-Year Perspective*. Springer, 1999.
- [BBB⁺14] Maurice Bruynooghe, Hendrik Blockeel, Bart Bogaerts, Broes De Cat, Stef De Pooter, Joachim Jansen, Anthony Labarre, Jan Ramon, Marc Denecker, and Sicco Verwer. Predicate logic as a modeling language: Modeling and solving some machine learning and data mining problems with IDP3. *Theory and Practice of Logic Programming*, pages 1–35, 2014.
- [Bec05a] Moritz Y. Becker. Cassandra: Flexible trust management and its application to electronic health records. PhD dissertation, Technical Report UCAM-CL-TR-648, Computer Laboratory, University of Cambridge, 2005.
- [Bec05b] Moritz Y. Becker. A formal security policy for an NHS electronic health record service. Technical Report UCAM-CL-TR-628, Computer Laboratory, University of Cambridge, 2005.
- [Ber13] Bloom Programming Language. <http://www.bloom-lang.net>, 2013. Latest release April 23, 2013. Accessed January 14, 2017.
- [BF06] Steve Barker and Maribel Fernández. Term rewriting for access control. In *Data and applications security XX*, pages 179–193. Springer, 2006.
- [BFL96] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, pages 164–173, 1996.
- [BG05] Marcello Balduccini and Michael Gelfond. Model-based reasoning for complex flight systems. In *Proceedings of the 5th AIAA Conference on Aviation, Technology, Integration, and Operations*, 2005.
- [BK94] Anthony J. Bonner and Michael Kifer. An overview of transaction logic. *Theoretical Computer Science*, 133(2):205–265, 1994.
- [BKB14] Reza Basseda, Michael Kifer, and Anthony J Bonner. Planning with transaction logic. In *Proceedings of the 8th International Conference on Web Reasoning and Rule Systems*, pages 29–44. Springer, 2014.
- [BKGD⁺12] Anthony Burton, Robert Kowalski, Marta Gacic-Dobo, Rouslan Karimov, and David Brown. A formal representation of the WHO and UNICEF estimates of national immunization coverage: A computational logic approach. *PLOS ONE*, Oct. 2012.
- [BLV04] Steve Barker, Michael Leuschel, and Mauricio Varea. Efficient and flexible access control via logic program specialisation. In *Proceedings of the 2004 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 190–199, 2004.
- [Bon10] Piero A. Bonatti. Datalog for security, privacy and trust. In *Proceedings of the 1st International Conference on Datalog Reloaded*, pages 21–36. Springer, 2010.
- [BRS12] Moritz Y Becker, Alessandra Russo, and Nik Sultana. Foundations of logic-based trust management. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, pages 161–175. IEEE CS Press, 2012.
- [BS04] Moritz Y. Becker and Peter Sewell. Cassandra: Flexible trust management, applied to electronic health records. In *Proceedings of the 17th IEEE Computer Security Foundations Workshop*, pages 139–154. IEEE CS Press, 2004.
- [BS09a] Jordan Bell and Brett Stevens. A survey of known results and research areas for n-queens. *Discrete Mathematics*, 309(1):1–31, 2009.

- [BS09b] Martin Bravenboer and Yannis Smaragdakis. Strictly declarative specification of sophisticated points-to analyses. In *Proceedings of the 24rd ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications*, pages 243–262, 2009.
- [CEFP14] Alessandro Cimatti, Stefan Edelkamp, Maria Fox, and Erion Plaku. Dagstuhl Seminar 14482: Automated Planning and Model Checking. http://www.dagstuhl.de/no_cache/en/program/calendar/semhp/?semnr=14482, Nov. 23–28, 2014. Accessed June 6, 2015.
- [CGL94] Edmund M. Clarke, Orna Grumberg, and David E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems*, 16(5):1512–1542, 1994.
- [CGP99] Edmund M. Clarke, Jr., Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press, 1999.
- [CGT90] Stefano Ceri, Georg Gottlob, and Letizia Tanca. *Logic Programming and Databases*. Springer, 1990.
- [CS95] P. Ciancarini and Leon Sterling. Report on the Workshop: Applications of Logic Programming in Software Engineering. *The Knowledge Engineering Review*, 10(01):97–100, 1995.
- [DAA13] Carlos Viegas Damásio, Anastasia Analyti, and Grigoris Antoniou. Justifications for logic programming. In *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning*, pages 530–542. Springer, 2013.
- [DEGV01] Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3):374–425, 2001.
- [DeT02] John DeTreville. Binder, a logic-based security language. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 105–113. IEEE CS Press, 2002.
- [dMLW03] Oege de Moor, David Lacey, and Eric Van Wyk. Universal regular path queries. *Higher-Order and Symbolic Computation*, 16(1–2):15–35, 2003.
- [DNST05] Bart Demoen, Phuong-Lan Nguyen, Tom Schrijvers, and Remko Troncon. The first 10 Prolog programming contests. <http://dtai.cs.kuleuven.be/ppcbook/>, 2005. Accessed May 20, 2015.
- [DP10] Agostino Dovier and Enrico Pontelli, editors. *A 25-year Perspective on Logic Programming: Achievements of the Italian Association for Logic Programming, GULP*. Springer, 2010.
- [DPT10] Alessandro Dal Palù and Paolo Torroni. 25 years of applications of logic programming in Italy. In *A 25-Year Perspective on Logic Programming*, pages 300–328. Springer, 2010.
- [DRKT07] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. ProbLog: A probabilistic Prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2468–2473. Morgan Kaufman, 2007.
- [DT08] Marc Denecker and Eugenia Ternovska. A logic of nonmonotone inductive definitions. *ACM Transactions on Computational Logic*, 9(2):14, 2008.
- [Edm15] Doug Edmunds. Learning constraint logic programming—finite domains with logic puzzles. <http://brownbuffalo.sourceforge.net/>, 2015. Accessed May 20, 2015.
- [EFL⁺99] C. Ellison, B. Frantz, B. Lampson, R. L. Rivest, B. Thomas, and T. Ylonen. RFC 2693: SPKI Certificate Theory. <http://www.ietf.org/rfc/rfc2693.txt>, Sept. 1999. Accessed June 4, 2015.
- [EMP14] Stefan Edelkamp, Daniele Magazzeni, and Erion Plaku. Workshop on Model Checking and Automated Planning (MOCHAP’14). http://icaps14.icaps-conference.org/workshops_tutorials/mochap.html, Portsmouth, NH, June 23, 2014. Accessed June 6, 2015.
- [Fit02] Melvin Fitting. Fixpoint semantics for logic programming: A survey. *Theoretical Computer Science*, 278(1):25–51, 2002.
- [FK92] D. Ferraiolo and R. Kuhn. Role-based access control. In *Proceedings of the 15th NIST-NSA National Computer Security Conference*, pages 554–563, Blatimore, Maryland, 1992. <http://arxiv.org/abs/0903.2171>.

- [FSG⁺01] David F. Ferraiolo, Ravi Sandhu, Serban Gavrilă, D. Richard Kuhn, and Ramaswamy Chandramouli. Proposed NIST standard for role-based access control. *ACM Transactions on Information and Systems Security*, 4(3):224–274, 2001.
- [GAK12] Todd J. Green, Molham Aref, and Grigoris Karvounarakis. LogicBlox, platform and language: A tutorial. In *Proceedings of the 2nd International Conference on Datalog in Academia and Industry*, Datalog 2.0, pages 1–8. Springer, 2012.
- [GAS10] Ana Sofia Gomes, José Júlio Alferes, and Terrance Swift. Implementing query answering for hybrid MKNF knowledge bases. In *Proceedings of the 12th International Conference on Practical Aspects of Declarative Languages*, pages 25–39. Springer, 2010.
- [GB98] A. Gavrilă and J. Barkley. Formal specification for RBAC user/role and role relationship management. In *Proceedings of the 3rd ACM Workshop on Role Based Access Control*, pages 81–90, 1998.
- [GBF⁺15] Benjamin Grosz, Janine Bloomfield, Paul Fodor, Michael Kifer, Isaac Grosz, Miguel Calejo, and Theresa Swift. Automated decision support for financial regulatory/policy compliance, using textual rulelog. In *Proceedings of the RuleML 2015 Challenge, the Special Track on Rule-based Recommender Systems for the Web of Data, the Special Industry Track and the RuleML 2015 Doctoral Consortium*, 2015. <http://ceur-ws.org/Vol-1417/>.
- [GILR09] Giovanni Grasso, Salvatore Iiritano, Nicola Leone, and Francesco Ricca. Some DLV applications for knowledge management. In E. Erdem, F. Lin, and T. Schaub, editors, *Proceedings of the 10th International Conference on Logic Programming and Nonmonotonic Reasoning*, pages 591–597. Springer, 2009.
- [GKKS12] M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool, 2012.
- [GL88] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In *Proceedings of the 5th International Conference and Symposium on Logic Programming*, pages 1070–1080. MIT Press, 1988.
- [GLMR11] Giovanni Grasso, Nicola Leone, Marco Manna, and Francesco Ricca. ASP at work: Spin-off and applications of the DLV system. In *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning—Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday*, pages 432–451. Springer, 2011.
- [GLR13] Giovanni Grasso, Nicola Leone, and Francesco Ricca. Answer set programming: Language, applications and development tools. In *Proceedings of the 7th International Conference on Web Reasoning and Rule Systems*, pages 19–34. Springer, 2013.
- [GM01] Harald Ganzinger and David A. McAllester. A new meta-complexity theorem for bottom-up logic programs. In *Proceedings of the 1st International Joint Conference on Automated Reasoning*, pages 514–528. Springer, 2001.
- [GMR17] Martin Gebser, Marco Maratea, and Francesco Ricca. The sixth answer set programming competition. *J. Artif. Intell. Res.*, 60:41–95, 2017.
- [Gro05] TU Wien Knowledge-Based Systems Group. WP5 report: Model applications and proofs-of-concept. <http://www.kr.tuwien.ac.at/research/projects/WASP/report.pdf>, Aug. 2005. Accessed May 20, 2015.
- [GS00] Tyrone Grandison and Morris Sloman. A survey of trust in Internet applications. *IEEE Communications Surveys and Tutorials*, 3(4):2–16, 2000.
- [HDGD10] P. Hou, B. De Cat, and M. Denecker. FO(FD): Extending classical logic with rule-based fixpoint definitions. *Theory and Practice of Logic Programming*, 10(4-6):581–596, 2010.
- [Het15] Werner Hett. Prolog Site—Prolog Problems. <http://sites.google.com/site/prologsite/prolog-problems/>, 2015. Accessed May 28, 2015.
- [HGL11] Shan Shan Huang, Todd Jeffrey Green, and Boon Thau Loo. Datalog and emerging applications: An interactive tutorial. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 1213–1216, 2011.

- [Hin01] Michael Hind. Pointer analysis: Haven't we solved this problem yet? In *Proceedings of the 2001 ACM SIGPLAN-SIGSOFT workshop on Program Analysis for Software Tools and Engineering*, pages 54–61, 2001.
- [HT01a] Nevin Heintze and Olivier Tardieu. Demand-driven pointer analysis. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation*, pages 24–34, 2001.
- [HT01b] Nevin Heintze and Olivier Tardieu. Ultra-fast aliasing analysis using CLA: A million lines of C code in a second. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation*, pages 254–263, 2001.
- [HTL07] Katia Hristova, K. Tuncay Tekle, and Yanhong A. Liu. Efficient trust management policy analysis from rules. In *Proceedings of the 9th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, pages 211–220, 2007.
- [HVM06] Elnar Hajiyev, Mathieu Verbaere, and Oege De Moor. CodeQuest: Scalable source code queries with Datalog. In *Proceedings of the 20th European Conference on Object-Oriented Programming*, pages 2–27. Springer, 2006.
- [Ioa96] Yannis E Ioannidis. Query optimization. *ACM Computing Surveys*, 28(1):121–123, Mar. 1996.
- [Jim01] Trevor Jim. SD3: A trust management system with certified evaluation. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 106–115. IEEE CS Press, 2001.
- [JM94] Joxan Jaffar and Michael J. Maher. Constraint logic programming: A survey. *Journal of Logic Programming*, 19:503–581, 1994.
- [KBL06] Michael Kifer, Arthur Bernstein, and Philip M. Lewis. *Database Systems: An Application Oriented Approach, Complete Version*. Addison-Wesley, 2nd edition, 2006.
- [Kje15] Hakan Kjellerstrand. My Picat page. <http://www.hakank.org/picat/>, 2015. Accessed May 29, 2015.
- [KLW95] Michael Kifer, Georg Lausen, and James Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42(4):741–843, 1995.
- [Kow14] Robert Kowalski. Logic programming. In Dov M. Gabbay, Jörg H. Siekmann, and John Woods, editors, *Computational Logic*, volume 9 of *Handbook of the History of Logic*, pages 523–569. Elsevier, 2014.
- [KS15] Robert Kowalski and Fariba Sadri. Reactive computing as model generation. *New Generation Computing*, 33(1):33–67, 2015.
- [KYWZ14] Michael Kifer, Guizhen Yang, Hui Wan, and Chang Zhao. *Flora-2: User's Manual Version 1.0*. Stony Brook University, July 2014. <http://flora.sourceforge.net/>. Accessed June 6, 2015.
- [LB08] Michael Leuschel and Michael Butler. ProB: An automated analysis toolset for the B method. *International Journal on Software Tools for Technology Transfer*, 10(2):185–203, 2008.
- [LBSL16] Yanhong A. Liu, Jon Brandvein, Scott D. Stoller, and Bo Lin. Demand-driven incremental object queries. In *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming*, pages 228–241. ACM Press, 2016.
- [LCG⁺09] Boon Thau Loo, Tyson Condie, Minos Garofalakis, David E. Gay, Joseph M. Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. Declarative networking. *Communications of the ACM*, 52:87–95, 2009.
- [Ley02] Michael Ley. The DBLP computer science bibliography: Evolution, research issues, perspectives. In *Proceedings of the 9th International Symposium on String Processing and Information Retrieval*, pages 1–10. Springer, 2002.
- [LF11] Adam Lally and Paul Fodor. Natural language processing with Prolog in the IBM Watson system. *Association for Logic Programming (ALP) Issue, Featured Articles*, Mar. 31 2011. Accessed April 23, 2015.
- [LHM84] Carl E. Landwehr, Constance L. Heitmeyer, and John McLean. A security model for military message systems. *ACM Transactions on Computer Systems*, 2(3):198–222, 1984.

- [LM03] Ninghui Li and John C. Mitchell. Datalog with constraints: A foundation for trust management languages. In *Proceedings of the 5th International Symposium on Practical Aspects of Declarative Languages*, pages 58–73. Springer, 2003.
- [LMW02] Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust-management framework. In *IEEE Symposium on Security and Privacy*, pages 114–130, 2002.
- [Log15a] LogicBlox. Assortment planning and management. <http://www.logicblox.com/solution-four.html>, 2015. Accessed May 18, 2015.
- [Log15b] LogicBlox. Solutions. <http://www.logicblox.com/solutions.html>, 2015. Accessed May 18, 2015.
- [LPF⁺06] Nicola Leone, Gerald Pfeifer, Wolfgang Faber, Thomas Eiter, Georg Gottlob, Simona Perri, and Francesco Scarcello. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic*, 7(3):499–562, July 2006.
- [LPM⁺12] Adam Lally, John M. Prager, Michael C. McCord, Branimir K. Boguraev, Siddharth Patwardhan, James Fan, Paul Fodor, and Jennifer Chu-Carroll. Question analysis: How Watson reads a clue. *IBM Journal of Research and Development*, 56(3/4):2:1–2:13, 2012.
- [LR15] Nicola Leone and Francesco Ricca. Answer Set Programming: A tour from the basics to advanced development tools and industrial applications. In *Proceedings of the 11th International Summer School on Reasoning Web*, pages 308–326. Springer, 2015.
- [LRY⁺04] Yanhong A. Liu, Tom Rothamel, Fuxiang Yu, Scott Stoller, and Nanjun Hu. Parametric regular path queries. In *Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation*, pages 219–230, 2004.
- [LS06] Yanhong A. Liu and Scott D. Stoller. Querying complex graphs. In *Proceedings of the 8th International Symposium on Practical Aspects of Declarative Languages*, pages 199–214. Springer, 2006.
- [LS07] Yanhong A. Liu and Scott D. Stoller. Role-based access control: A corrected and simplified specification. In *Department of Defense Sponsored Information Security Research: New Methods for Protecting Against Cyber Threats*, pages 425–439. Wiley, 2007.
- [LS09] Yanhong A. Liu and Scott D. Stoller. From Datalog rules to efficient programs with time and space guarantees. *ACM Transactions on Programming Languages and Systems*, 31(6):1–38, 2009.
- [LS18] Yanhong A. Liu and Scott D. Stoller. Founded semantics and constraint semantics of logic rules. In *Symposium on Logical Foundations of Computer Science*, Lecture Notes in Computer Science. Springer, Jan. 2018.
- [LWG⁺06] Yanhong A. Liu, Chen Wang, Michael Gorbovitski, Tom Rothamel, Yongxi Cheng, Yingchao Zhao, and Jing Zhang. Core role-based access control: Efficient implementations by transformations. In *Proceedings of the ACM SIGPLAN 2006 Workshop on Partial Evaluation and Program Manipulation*, pages 112–120, 2006.
- [Mal15] Mihaela Malita. Logic puzzles in Prolog. http://www.anselm.edu/internet/compsci/faculty_staff/mmalita/HOMEPAGE/logic/index.html, 2015. Accessed May 28, 2015.
- [McA99] David A. McAllester. On the complexity analysis of static analyses. In *Proceedings of the 6th International Static Analysis Symposium*, pages 312–329. Springer, 1999.
- [MSZ14] Jack Minker, Dietmar Seipel, and Carlo Zaniolo. Logic and databases: History of deductive databases. In D. Gabbay, J. Siekmann, and J. Woods, editors, *Handbook of Computational Logic*, chapter 17, pages 571–628. North-Holland, 2014.
- [NBG⁺01] Monica Nogueira, Marcello Balduccini, Michael Gelfond, Richard Watson, and Matthew Barry. An A-Prolog decision support system for the Space Shuttle. In *Practical Aspects of Declarative Languages*, pages 169–183. Springer, 2001.
- [Prz94] Teodor C. Przymusiński. Well-founded and stationary models of logic programs. *Annals of Mathematics and Artificial Intelligence*, 12(3):141–187, 1994.

- [PS02] Fernando C.N. Pereira and Stuart M Shieber. *Prolog and Natural-Language Analysis*. Microtome Publishing, 2002. Revision of October 5, 2005.
- [RGA⁺12] Francesco Ricca, Giovanni Grasso, Mario Alviano, Marco Manna, Vincenzino Lio, Salvatore Iritano, and Nicola Leone. Team-building with answer set programming in the Gioia-Tauro Seaport. *Theory and Practice of Logic Programming*, 12(3):361–381, 2012.
- [RGS⁺09] Francesco Ricca, Lorenzo Gallucci, Roman Schindlauer, Tina Dell’Armi, Giovanni Grasso, and Nicola Leone. OntoDLV: An ASP-based system for enterprise ontologies. *Journal of logic and computation*, 19(4):643–670, 2009.
- [RK05] Sini Ruohomaa and Lea Kutvonen. Trust management survey. In *Proceedings of the Third international conference on Trust Management*, pages 77–92. Springer, 2005.
- [RL07] Tom Rothamel and Yanhong A. Liu. Efficient implementation of tuple pattern based retrieval. In *Proceedings of the ACM SIGPLAN 2007 Workshop on Partial Evaluation and Program Manipulation*, pages 81–90, 2007.
- [Rot93] Al Roth. The practical application of Prolog. *AI Expert*, 8:24–24, 1993. In Dr. Dobb’s, <http://www.drdobbs.com/parallel/the-practical-application-of-prolog/184405220>, Dec.10, 2002. Accessed June 6, 2015.
- [RRR00] Abhik Roychoudhury, CR Ramakrishnan, and IV Ramakrishnan. Justifying proofs using memo tables. In *Proceedings of the 2nd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, pages 178–189, 2000.
- [RRS⁺00] C.R. Ramakrishnan, I.V. Ramakrishnan, Scott A. Smolka, Yifei Dong, Xiaoqun Du, Abhik Roychoudhury, and V.N. Venkatakrishnan. XMC: A logic-programming-based verification toolset. In *Proceedings of the 12th International Conference on Computer Aided Verification*, pages 576–580. Springer, 2000.
- [RU95] Raghu Ramakrishnan and Jeffrey D Ullman. A survey of deductive database systems. *Journal of Logic Programming*, 23(2):125–149, 1995.
- [SB15] Yannis Smaragdakis and George Balatsouras. Pointer analysis. *Foundations and Trends in Programming Languages*, 2(1):1–69, 2015.
- [SBK13] Nik Sultana, Moritz Y. Becker, and Markulf Kohlweiss. Selective disclosure in Datalog-based trust management. In *Proceedings of the 9th International Workshop on Security and Trust Management*, pages 160–175. Springer, 2013.
- [SCD⁺13] Manu Sridharan, Satish Chandra, Julian Dolby, Stephen J Fink, and Eran Yahav. Alias analysis for object-oriented programs. In *Aliasing in Object-Oriented Programming: Types, Analysis and Verification*, pages 196–232. Springer, 2013.
- [Sch11] Torsten Schaub. Collection on Answer Set Programming (ASP) and more. <http://www.cs.uni-potsdam.de/~torsten/asp/>, Mar. 2011. Accessed May 18, 2015.
- [Sch14] Torsten Schaub. Answer set solving in practice. <http://www.cs.uni-potsdam.de/~torsten/Potassco/Slides/asp.pdf>, Dec. 23, 2014. Accessed May 20, 2015.
- [Sem12] Semafora. Semantic infrastructure: OntoBroker. <http://www.semafora-systems.com/en/products/ontobroker/>, 2012. Accessed May 18, 2015.
- [SK97] Taisuke Sato and Yoshitaka Kameya. PRISM: A language for symbolic-statistical modeling. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence-Volume 2*, pages 1330–1335. Morgan Kaufman, 1997.
- [SR05] Diptikalyan Saha and C. R. Ramakrishnan. Incremental and demand-driven points-to analysis using logic programming. In *Proceedings of the 7th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, pages 117–128, 2005.
- [SW12] Terrance Swift and David S Warren. XSB: Extending Prolog with tabled logic programming. *Theory and Practice of Logic Programming*, 12(1-2):157–187, 2012.
- [SW⁺14] Terrance Swift, David S. Warren, et al. *The XSB System Version 3.5.x*, June 2014. <http://xsb.sourceforge.net>. Accessed June 6, 2015.

- [SYGR11] Scott D Stoller, Ping Yang, Mikhail I Gofman, and CR Ramakrishnan. Symbolic reachability analysis for parameterized administrative role-based access control. *Computers & Security*, 30(2):148–164, 2011.
- [TGL10] K. Tuncay Tekle, Michael Gorbovitski, and Yanhong A. Liu. Graph queries through Datalog optimizations. In *Proceedings of the 12th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming*, pages 25–34, 2010.
- [TL11] K. Tuncay Tekle and Yanhong A. Liu. More efficient Datalog queries: Subsumptive tabling beats magic sets. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, pages 661–672, 2011.
- [VG93] Allen Van Gelder. The alternating fixpoint of logic programs with negation. *Journal of Computer and System Sciences*, 47(1):185–221, 1993.
- [VRS91] Allen Van Gelder, Kenneth Ross, and John S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.
- [WACL05] John Whaley, Dzintars Avots, Michael Carbin, and Monica S Lam. Using Datalog with binary decision diagrams for program analysis. In *Programming Languages and Systems*, pages 97–118. Springer, 2005.
- [Wil02] Dan E. Willard. An algorithm for handling many relational calculus queries efficiently. *Journal of Computer and System Sciences*, 65:295–331, 2002.
- [WL17] David S. Warren and Yanhong A. Liu. AppLP: A dialogue on applications of logic programming. *Computing Research Repository*, arXiv:1704.02375 [cs.PL], Apr. 2017.
- [XSB15] XSB. Case Studies. <http://www.xsb.com/case-studies>, 2015. Accessed May 18, 2015.