Noname manuscript No.
(will be inserted by the editor)

Reciprocal influences between proof theory and logic programming

the date of receipt and acceptance should be inserted later

Abstract The topics of structural proof theory and logic programming have influenced each other for more than three decades. Proof theory has contributed the notion of sequent calculus, linear logic, and higher-order quantification. Logic programming has forced an examination of what are synthetic inference rules and what is a good logic-based approach to the treatment of bindings. As a result, proof theory has responded by developing an approach to proof search that is based on focused proof systems in which introduction rules are organized into two alternating phases of rule application. Since the logic programming community can generate many examples and many design goals (e.g., modularity of specifications, higher-order programming, etc.), the close connections with proof theory have helped to keep the latter community relevant to the large world of computational logic.

Keywords Structural proof theory · logic programming · computational logic · history of programming languages

1 Introduction

Both symbolic logic and the theory of proof have been successfully used in the foundations of mathematics. For example, Gentzen's early work on the sequent calculus was used to show the consistency of classical and intuitionistic logic and arithmetic. The last several decades have clearly demonstrated that logic has a significant and continuing impact on computer science, possibly rivaling its impact on mathematics. For example, there are major journals that cover the general topic of computational logic—the ACM Transactions on Computational Logic, Logical Methods in Computer Science, the Journal on Automated Reasoning, and the Journal of Logic and Computation—to name a few. Similarly, there are several major conferences (e.g., CADE, CSL, FSCD,

LICS, IJCAR) that similarly address various uses of logic in computational settings. This topic also has its own "unreasonable effectiveness" paper, namely "On the Unusual Effectiveness of Logic in Computer Science" [33].

As computer science moves forward, researchers and practitioners often design ad hoc programs, protocols, data structures, etc. Usually, the first pressures applied to the designers of such systems is the need to support efficient implementation or other, short-term goals (e.g., interoperability with existing code and hardware). Other pressures often appear later: e.g., it might be important that the code be modular; that it has good properties that can be determined before its execution (such as type checking and various forms of static analysis); or that correctness issues are important and systems designs must facilitate formal reasoning.

While short-term fixes can always be realized, poor designs can lead to long-term costs. On this point, it is useful to be reminded of the following, often-cited quote.

> "Beauty is the first test: there is no permanent place in the world for ugly mathematics." – G. H. Hardy, *A Mathematician's Apology* [34]

What a computer scientist can take from this quote is that an ugly computer system (a "kludge"), even one that might be working, may have no permanent place in the world: all the various pressure imposed on computing will likely force ugly systems to be replaced by better designs. For example, to expect that a program can be compiled so that it works on a single processor as well as on multiple processors, deep understanding of the semantics and formal properties of a programming language is required: quickly hacked programs do not support deep understanding nor formal properties.

When looking to articulate and exploit deep principles in computing, researchers are often led to develop new frameworks. For example, finite state machines and context-free grammars have been developed and exploited to provide a strong foundation for parsing strings into structured data. When needing to deal with communications and shared resources in computer networks, process calculi, such as CSP [37] and CCS [75], have been developed, studied, and shaped into programming languages (e.g., the Occam programming language [12]). Occasionally, syntactic systems that are not traditionally considered logics are so well studied and found to be of such high quality that they can be used as frameworks for programming languages: the λ-calculus [8, 17] and the π-calculus [76, 77, 92] are two such examples.

In this paper, I will show how various features of some well studied logical systems directly influenced aspects of programming. At the same time, I will also provide some examples where attempts to deal with various needs of computing directly lead to new designs and results in logic. Logic is a challenging framework for computation: much can be gained by rising to that challenge to find logical principles behind computation.

I should make it clear before proceeding that I am a participant in the more than three-decade long story that I give in this paper: I am not a detached and objective historian. I have two goals in mind in telling this story. First, that

my experience and hindsight in the field will interest those who want specific examples of the mutual influence that has been exerted between the abstract and formal topics of proof theory and the concrete and practical topic of computer programming languages. Second, I want to show how a certain part of computer science can be attached to the foundations of formal proof that was pioneered by Hilbert, Gödel, and Gentzen: the foundation that they and many other provided in the first half of the 20th century have had significant and immediate impact on certain topics of computer science even today.

2 Logic and computation: the different uses of logic

Early in the 20th century, some logicians invented various computational systems, such as Turing machines, Church's λ-calculus, and Post correspondence systems, which were shown to all compute the same set of *recursive functions*. With the introduction of high-level programming languages, such as LISP, Pascal, Ada, and C, it was clear that any number of ad hoc computation systems could be designed to compute these same functions. Eventually, the large number of different programming languages were classified via the four paradigms of *imperative*, *object-oriented*, *functional*, and *logic* programming. The latter two can be viewed as an attempt to make less ad hoc computational systems by relying on aspects of symbolic logic. Unlike most programming languages, symbolic logic is a formal language that has well-defined semantics and which has been studied using model theory [96], category theory [51, 52], recursion theory [30, 45], and proof theory [28, 29]. As we now outline, logic has taken on a number of different roles when it is applied to computation.

The earliest and most popular use of logic in computer science views computation as something that happens independently of logic: e.g., registers change, tokens move in a Petri net, messages are buffered and retrieved, and a tape head advances along a tape. Logics (often modal or temporal logics) are used to make statements *about* such computations. Model checkers and Hoare proof systems employ this *computation-as-model* approach.

Another use of logic is to provide specification and programming languages whose syntax and semantics are tied directly to logic. The *computation-as-deduction* approach to programming languages takes as its computational elements objects from logic, namely, types, terms, formulas, and proofs. Thus, instead of ad hoc devices to encode computation (e.g., the tape of tokens in a Turing machine), these elements of logic would have, presumably, canonical meaning which can be found in mathematical texts instead of the specific technology of particular compilers and interpreters. Also, programs that rely heavily on logic-based formalism might be able to exploit the rich metatheory of logic to help in proving properties of specific programs and entire programming languages.

There are, however, two broad and strikingly different ways to apply the computation-as-deduction approach to modeling computation: these different

avenues rely on different roles of proof in the design and analysis of computation.

Proof normalization: Programs can be viewed as proofs of their types and computation is viewed as a process of proof normalization. Such normalization is generally based on either β-reduction using natural deduction proofs or on cut-elimination using sequent calculus proofs. This proof-as-program approach is often used to provide formal models of functional programming.

Proof search: Formulas can be used to encode both programs and goals (think to rules and queries in database theory). Sequents are used to encode the state of a computation and (cut-free) proof search is used to provide traces in computation: changes in sequents denote the dynamics of computation. Cut-elimination is not part of computation but can be used to reason about computation. This view of computation is used to provide a foundation to logic programming.

Although both of these frameworks put formal proofs at their core, the difference between these two approaches is a persistent one. For example, advances in the understanding of the proof theory of higher-order quantification and linear logic have resulted in different advances in both of these paradigms separately. No convergence of these two topics has been forced from any advances in our understanding of proof and their underlying logics.

The connections between functional programming and proof theory is well documented and celebrated in the literature as the *Curry-Howard Isomorphism*: see, for example, [84, 94]. The connection between logic programming and proof theory is less well documented and it is the focus of this article.

The field of proof theory covers many topics, including consistency proofs, ordinal inductions, reverse mathematics, proof mining, and proof complexity. Here, we focus instead on *structural proof theory*, a topic initiated by Gentzen's introduction of sequent calculus and natural deduction [28]. The sequent calculus is particularly appealing since Gentzen explicitly preferred it over natural deduction as a setting for unifying the concept of proof in both classical and intuitionistic logic. Later, Girard showed that the sequent calculus can naturally account for linear logic proofs as well [29]. As we shall document, this feature of the sequent calculus provides logic programming a natural framework in which proof-search could be described for much richer logics (first-order and higher-order versions of classical, intuitionistic, and linear logics) than that underlying Prolog. Another feature of sequent calculus is its support for *abstraction*: that is, it provides mechanisms for allowing some aspects of a program's specification to be hidden while other aspects are made explicit. In programming language terminology, such abstractions provide logic programming with modularity, abstract data types, and higher-order programming. The use of abstractions can significantly aid in establishing formal properties of programs [62].

3 A quick primer: terms and formulas of predicate logic

We shall assume that the reader has at least some familiarity with first-order predicate logic. In this section, we simply review a few concepts that will help to anchor our later discussions.

In order to define term and formula structures, we need to know which symbols denote predicates and function symbols and what is their arity. Many first-order logic systems (including most Prolog languages) only declare the arity of such symbols. For example, the constructors for natural numbers and lists of natural numbers can be written as

$$\{z/0,\ s/1,\ \mathit{nil}/0,\ \mathit{cons}/2,\ \mathit{append}/3\}$$

Thus, *cons* (the non-empty list constructor) takes two arguments while *append* (the relation between two lists and the result of appending them) takes three arguments. Some first-order logics are sorted: that is, there are primitive sorts, say, *nat* (for natural numbers) and *list* (for lists of natural numbers), and constructors are declared to take their arguments from certain sorts. For example, the declaration displayed above could be made more explicit using sorts such as

$$\{z : \langle\langle\rangle, \mathit{nat}\rangle,\ s : \langle\langle \mathit{nat}\rangle, \mathit{nat}\rangle,\ \mathit{nil} : \langle\langle\rangle, \mathit{list}\rangle, \mathit{cons} : \langle\langle \mathit{nat}, \mathit{list}\rangle, \mathit{list}\rangle\}$$

Above only term constructors are given declarations in which the first member of their associated tuple is the list of argument types it expects and the second member is the type of the object that the constructor builds. Predicates could be declared separately using the declaration $\{\mathit{append} : \langle \mathit{list}, \mathit{list}, \mathit{list}\rangle\}$ which associates a predicate with the list of argument types it expects.

So that we can also comment on higher-order logic and syntax later, we use the conventions introduced by Church's Simple Theory of Types [16]. In particular, a *type* is either a *primitive type* (these are introduced as we need them and correspond to primitive sorts) and an *arrow type* which is an expression of the form $\tau_1 \to \tau_2$. The arrow associates to the right: thus $\tau_1 \to \tau_2 \to \cdots \to \tau_n \to \tau_0$ reads as $\tau_1 \to (\tau_2 \to \cdots \to (\tau_n \to \tau_0)\cdots)$. A function symbol with the sort declaration $\langle\langle\tau_1, \ldots, \tau_n\rangle, \tau_0\rangle$ would correspond to the type $\tau_1 \to \tau_2 \to \cdots \to \tau_n \to \tau_0$. A predicate symbol with the sort declaration $\langle\tau_1, \ldots, \tau_n\rangle$ is encoded as the type $\tau_1 \to \tau_2 \to \cdots \to \tau_n \to o$, where we follow Church's convention to use the primitive type o to denote the (syntactic category of) formulas. Thus, the declarations above can be revised to be

$$\begin{array}{l}\{z : \mathit{nat},\ s : \mathit{nat} \to \mathit{nat},\ \mathit{nil} : \mathit{list},\ \mathit{cons} : \mathit{nat} \to \mathit{list} \to \mathit{list},\\ \quad \mathit{append} : \mathit{list} \to \mathit{list} \to \mathit{list} \to o\}\end{array}$$

While the arrow type constructor is natural for presenting first-order logic, its presence will also make it easy to generalize the syntax of terms and formulas to accommodate higher-order logic (in Section 8).

A *signature* is a set containing pairings of tokens with their declared type so that all tokens are declared to have at most one type. Informally, a Σ-term of type τ is a (closed) term all of whose tokens are taken from the

signature Σ and which respects the typing declarations. For example, if Σ' is the signature declared at the end of the previous paragraph, then $(s\ (s\ z))$ and $(cons\ (s\ z)\ (cons\ z\ nil))$ are valid Σ-terms of type *nat* and *list*, respectively. The Σ'-term

$$(append\ (cons\ z\ nil)\ (cons\ (s\ z)\ nil)\ (cons\ z\ (cons\ (s\ z)\ nil)))$$

has type o which means that it is also a formula. In Prolog syntax, the latter expression corresponds to the (more compact) `append([0], [1], [0,1])`, which in turn denotes the assertion that `[0,1]` is the result of appending the lists `[0]` and `[1]`. In general, we intend the token *append* to stand for the three-place relation such that $(append\ L\ K\ M)$ holds if and only if the concatenation of the list L with the list K is the list M (a formal definition for this predicate is given in the next section).

The terms described above are examples of *closed* terms in the sense that they contain no free variables. Let $\mathcal{X}$ be an infinite set of token-type pairs of the form $x : \tau$ where τ is restricted to a primitive type. Assume that the two signatures Σ (of constants) and $\mathcal{X}$ (of first-order variables) do not contain the same token: in that case, a term over the combined signature $\Sigma \cup \mathcal{X}$ are terms with possible free variables.

Predicates symbols are introduced as a means to collect together some terms and to yield an *atomic formula* (such as the assertion about appending lists above). Non-atomic formulas are created using the following propositional constants (along with their declared types): $\top : o$ (truth), $\bot : o$ (false), $\neg : o \to o$ (negation), $\vee : o \to o \to o$ (disjunction), $\wedge : o \to o \to o$ (conjunction), and $\supset$: $o \to o \to o$ (implication). The two quantifiers are parameterized by a type: $\forall_\tau x.B$ and $\exists_\tau x.B$ denote the universal and existential quantifiers (respectively) of the variable x of type τ within the formula B. If a quantifier is written without a subscript type expression, then that type is either unimportant or is easy to infer from its context.

4 Early foundations of logic programming

The logic programming paradigm had a beginning within the Artificial Intelligence community dating back to the 1960s and 1970s. We start our story here with the first systematic development of a proof procedure by Kowalski [49], which provided a (nondeterministic) procedural interpretation of logic that lines up well with the nearly simultaneous development of the first Prolog system by Colmerauer [18].

4.1 Declarative vs procedural programs

Central to the beginnings of Prolog was the question of how it might be possible to take *declarative* information about a desired computation and to have that information turned into an actual procedure or program. For example, consider

the simple problem of concatenating two lists to get a third list. A declarative treatment of concatenation might be simply to state the following two facts about it.

1. Concatenating an empty list on the front of a list L yields the list L.
2. If the result of concatenating list L to the front of list K is the list M, then the result of concatenating list $(cons\ X\ L)$ to the front of list K is the list $(cons\ X\ M)$ for any X (of type *nat*).

Of course, there are many other statements about concatenation that one could make (for example, that concatenation is associative). These two facts can be captured easily in first-order logic. Using the predicate symbol *append* introduced in Section 3, the above two facts about concatenation can be encoded as the two formulas

$$\forall L(append\ nil\ L\ L) \quad \text{and}$$

$$\forall X \forall L \forall K \forall M[(append\ L\ K\ M) \supset (append\ (cons\ X\ L)\ K\ (cons\ X\ M))]$$

(Here, the type of X is *nat* and of L, K, and M is *list*.) Following standard Prolog-inspired conventions, we shall write variables as tokens with an initial capital letter and we shall drop all quantifiers assuming that all variables are universally quantified around such formulas. Another convention used by Prolog is to reverse the direction of the implication and to use an ASCII approximate `:-` to a turnstile ($\vdash$). Following all these conventions and shift font, we have the following Prolog-style program declaration.[1]

```
append nil L L.
append (cons X L) K (cons X M) :- append L K M.
```

For a second example of a declarative specification written using Prolog syntax, Figure 1 contains a small graph along with the specification of both the adjacency relation of that graph and a specification of the notion of a path between two points in that graph. In the last line of that specification, another Prolog convention is used: the comma denotes conjunction. That last line can be read as: if there is a step from X to Z and a path from Z to Y then there is a path from X to Y. We have also assumed that the signature for these formulas contains the following items

$$\begin{array}{l} a : node,\ b : node,\ c : node,\ d : node, \\ step : node \to node \to o,\ path : node \to node \to o \end{array}$$

where *node* is a primitive type denoting nodes in the graph.

In general, the logical formulas that underlie the Prolog programming language are formulas generally referred to as *Horn clauses*. These formulas are of the form

$$\forall x_1 \ldots \forall x_n[(A_1 \wedge \cdots \wedge A_n) \supset A_0], \qquad (n, m \geq 0)$$

[1] We actually us the syntax of λProlog instead of Prolog, although for simple programs, they are rather similar.

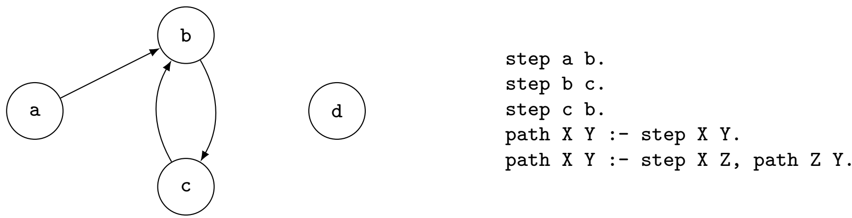

Fig. 1 A small graph on four nodes and a Prolog specification of it.

where the formulas $A_0, A_1, \ldots, A_n$ are atomic formulas all of whose free variables are in the set $\{x_1, \ldots, x_n\}$. If $n = 0$ then we do not write any universal quantifiers and if $m = 0$ then we do not write the implication. In classical logic, it is possible to convert all formulas to a logically equivalent formula (of essentially the same size) in which implication $\supset$ is not present and where all occurrences of negation $\neg$ are applied only to atomic formulas. Such formulas are in *negation normal form*. In particular, the negation normal form of the Horn clause above is

$$\forall x_1 \ldots \forall x_n [\neg A_1 \vee \cdots \vee \neg A_n \vee A_0].$$

If we let $\mathcal{P}$ be the set containing the five formulas displayed in Figure 1, it would seem natural to expect that provability from $\mathcal{P}$ and computing with this logic program might be related. For example, it is the case that $\mathcal{P}$ proves (in classical and intuitionistic logics) the atomic formula $(path\ a\ c)$ (i.e., that there is a path from node a to node c) and that the formula $(path\ a\ d)$ has no proof. While one might expect this connection to be rather immediate, the early history of Prolog obfuscated this connection with provability by describing logic programming computation as *refutation*, as we shall now illustrate.

4.2 Refutation and skolemization

In the late 1960's and early 1970's, the *resolution refutation* procedure of Robinson [90] was being applied in various areas of computational logic. For example, Green showed in [31] that resolution refutations could be used to provide answers within question-answering systems. Given the dominance of resolution, it was natural for Kowalski to have adopted it to provide a description of the operational behavior of Prolog.

On one hand, the choice of resolution was natural for this purpose since term unification was needed to describe Colmerauer's Prolog and since unification was built into the principle inference rule of resolution. On the other hand, this choice was unfortunate since it required turning what is most naturally considered a problem of searching for a *proof* into the problem of searching for a *refutation*. Since classical logic has an involutive negation, it is the case that proving A from $\mathcal{P}$ is equivalent to proving $\perp$ from $\mathcal{P} \cup \{A \supset \perp\}$: that

is, building a refutation of $\mathcal{P} \cup \{A \supset \perp\}$. Note that this latter step is not valid in intuitionistic logic: in general, resolution is not a sound procedure for intuitionistic logic (without significant modifications).

There seems to be only one reason why resolution and not proof dominated the early years of theorem proving in classical logic, and that was the use of *skolemization* to simplify quantifier structures in formulas. The process of skolemizing a first-order formula, say, B (in negation normal form) involves repeatedly replacing a subformula occurrence $\exists y.C(y)$ in B with $C(f(x_1, \ldots, x_n))$, where f is a new function symbol (an extension to the formula signature) and where $x_1, \ldots, x_n$ is the list of universally quantified variables of B that contain the occurrence $\exists y.C(y)$ in their scope. A Skolem normal form of B is then a formula that arises from repeatedly removing existential quantifiers in this manner until no occurrences of existential quantifiers remain. The main theorem that relates a formula B with a Skolem normal form of B is that they are equisatisfiable: that is, there is a model of B if and only if there is a model of a Skolem normal form of B. Since skolemization can introduce new constants (Skolem function symbols), the models of B are necessarily different from the models of a Skolem normal form of B. Thus, the stage is set for introducing refutations: in order to prove B is a theorem, we can show instead that $\neg B$ is unsatisfiable. This restatement is, of course, equivalent to showing that the skolemized form of $\neg B$ is unsatisfiable. It is this latter property that the resolution refutation framework is designed to demonstrate.

There are several reasons why involving skolemization and refutation into the theory of logic programming was not a good idea, at least in hindsight.

1. Horn clauses do not have quantifier alternations and, hence, skolemization is not a needed processing step.
2. Skolemization is not a sound process in intuitionistic logic. As we shall see, intuitionistic logic plays an important part in the future of logic programming.
3. Skolemization is a more complex and problematic process in the setting of higher-order quantification. For example, if one is not careful, Skolem functions can give rise to uses of the Axiom of Choice even for situations (such as logic programming) where one does not intend for the Axiom of Choice to be a relevant logical feature. Also, the result of a higher-order instantiation of a formula in Skolem normal form may result in a formula that is no longer Skolem normal (something that cannot happen in the first-order setting).

One benefit of using skolemization is that it makes the unification of terms a simple process: when eigenvariables are allowed in term structures, unification becomes a more complex operation. While modifications of unification to handle eigenvariables directly are known [64], such modifications arrived later and largely as a need to implement proof search mechanisms for intuitionistic logic. Similarly, various techniques, such as *deskolemization* [60], are available that allow Skolem terms and functions within a proof-theoretic setting: for

example, it is possible to use Skolem terms as particular naming schemes for eigenvariables within sequent calculus proofs.

4.3 SLD-resolution

There were two main ingredients in resolution refutations. The first ingredient is *clauses*, which are formulas of the form

$$\forall x_1 \ldots, \forall n_n[L_1 \vee \ldots \vee L_m], \qquad (n, m \geq 0)$$

where $x_1, \ldots, x_n$ is a (possibly empty) list of first-order variables and $L_1, \ldots, L_m$ is a (possibly empty) list of *literals* (atomic formulas or their negation). From what we noted above, Horn clauses can be seen as clauses in which exactly one literal is an atomic formula (instead of the negation of an atomic formula). In general, however, a clause can have any mixture of atomic formulas and negated atomic formulas.

The second ingredient is inference rules that take clauses as their premises and conclusion. The only one of these rules that interest us here is the so-called *resolution rule* which can be written as

$$\frac{\forall x_1 \ldots, \forall n_n[L \vee M] \qquad \forall y_1 \ldots, \forall y_m[\neg K \vee N]}{\forall z_1 \ldots, \forall z_p[\theta M \vee \theta N]}\ \theta = mgu(L, K).$$

Here, L and K are atomic formulas, M and N are (possibly empty) disjunctions of literals, and the proviso for this rule is that L and K are unifiable and that θ is set to the *most general unifier* L and K. A *resolution refutation* of the set of clauses $\{C_1, \ldots, C_q\}$ is a tree of such inference rules (plus another rule called *factoring*) in which the leaves come from the set of clauses and the root is the empty clause. When such a tree exists, the fact that the empty disjunction is clearly unsatisfiable can then be transfer to the collection of clauses in its leaves.

The resolution rule is rather remote from Gentzen's rules for sequent calculus. While Gentzen's introduction rules process exactly one logical connective per rule, the resolution rule above will deal with $n+m+p$ universal quantifiers along with a number of disjunctions. Furthermore, the operation of unification is not contained in sequent calculi presentations (although the implementation of theorem provers based on the sequent calculus often use unification).

Kowalski and Kuehner developed a specialized form of resolution based on *linear resolution with selection function* (SL-resolution) [50]. When this is variant of resolution is applied to Horn clauses, it was called SLD-resolution (D for definite) [7]. In this setting, attempting to prove the conjunctive goal $A_1 \wedge \ldots \wedge A_n$ from the Horn clauses in $\mathcal{P}$ results in attempting to refute the clauses in $\mathcal{P}$ together with the clause $\neg A_1 \vee \ldots \vee \neg A_n$: this latter clause is distinguished in that the literals it contain are all negated atoms. In this setting, SLD-resolution is essentially the restriction of resolution so that one of the clauses being used in the premise of a resolution is always the most recently produced such distinguished clause. This greatly restricted version of

resolution could be seen as forming the basis of the engine used in Prolog. Effective implementations of SLD-resolution were developed, with the most popular one based on the *Warren abstract machine* [1, 98].

Several variations on Horn clauses have been considered: these include *disjunctive logic programs* [55, 78] and *constraint logic programs* [42] (in which the primitive of equality is generalized to be a richer relations on terms (such as greater-than, non-equal-to): such constraints do not generally have most general solutions so one should not choose to solve them immediately but rather delay them until additional constraints are discovered. Most of these extensions were limited to features that could either be seen as retaining the basic characteristics of SLD-resolution or which could be compiled into the Warren Abstract Machine. While some extensions, such as HiLoG [15], proved useful in some circles, they often exerted no influence on the topics of logic and proof theory.

There are, however, many downsides of using resolution as the core explanation of how logic programming languages should work.

- Refuting is an odd choice in a setting where proving seems more natural.
- In order to present formulas as Horn clauses, one may need to transform a formula into its conjunction normal form, and this can cause an exponential increase in formula sizes or require the introduction of new predicate constants in order to keep that size from exploding.
- First-order unification maintains the normal form of clauses while this is not the case with higher-order quantification since predicate substitutions can transform a formula in normal form into one that is not in normal form. This particular problem could be addressed by re-normalizing after predicate substitutions [6, 41].
- More importantly, resolution does not naturally fit with intuitionistic and linear logics although it is possible to develop them based on the structure of sequent calculus proofs [24, 95].

These limitations with resolution refutations were then limitations to the designs of new logic programming languages. At roughly the same time as this framework was being designed for logic programming, researchers in functional programming languages were embracing many features of computational logic and proof theory that go well beyond the theory of first-order Horn clauses. In particular, higher-order programming, intuitionistic-logic based typing, and linear logic where all being considered as central and powerful themes in the design of modern functional programming languages. Guided by the Curry-Howard Isomorphism, the proof theory of higher-order intuitionistic logic helped to guide the design of a number of functional programming and reasoning systems [19, 56] and linear logic was seen as offering additional new features [97].

5 Proof theory characterization of Horn clauses

Gentzen's sequent calculus actually provides a natural setting for describing the operational behavior of proof search. Instead of building a refutation, one could instead attempt a proof. When attempting the proof of a goal G from a set of program clauses $\mathcal{P}$, we can consider the problem of building a Gentzen style proof system with the sequent $\mathcal{P} \vdash G$. For example, let $\mathcal{P}$ be a set of Horn clauses and let one of them be

$$\forall x_1 \ldots \forall x_m [(A_1 \wedge \cdots \wedge A_n) \supset A_0].$$

The following *backchaining* rule of inference is then *admissible* in Gentzen's LK calculus

$$\frac{\mathcal{P} \vdash \theta A_1 \quad \cdots \quad \mathcal{P} \vdash \theta A_n}{\mathcal{P} \vdash A}\ BC,$$

where it is the case that θ is a substitution such that $\Theta A_0 = A$. Admissibility of this rule is easy to see since it is the combination of one occurrence of a contraction, n occurrences of $\forall$-left introduction, one occurrence of $\supset$-left, and one occurrence of the initial rule. A stronger statement is also possible: this is the only inference rule that is required. That is, $\mathcal{P} \vdash A$ is provable in classical logic implies that there is a proof of that sequent in which only instances of the BC rule are needed.

Let A be a syntactic variable that ranges over first-order atomic formulas. Let $\mathcal{G}_1$ and $\mathcal{D}_1$ be the sets of all first-order G- and D-formulas defined inductively by the following rules:

$$G := \top \mid A \mid G_1 \wedge G_2 \mid G_1 \vee G_2 \mid \exists x\, G,$$
$$D := A \mid G \supset A \mid D_1 \wedge D_2 \mid \forall x\, D.$$

For the rest of this paper, the formulas of $\mathcal{D}_1$ are called *first-order Horn clauses.*

For the reader familiar with Church's treatment of higher-order logic, we define also a higher-order generalization to first-order Horn clauses. Let $\mathcal{H}_1$ be the set of all λ-normal terms that do not contain occurrences of the logical constants $\supset, \forall$, and $\perp$. Let A and A_r be syntactic variables denoting, respectively, atomic formulas and rigid atomic formulas (atomic formulas with a constant as its head symbol) in $\mathcal{H}_1$. Let $\mathcal{G}_2$ and $\mathcal{D}_2$ be the sets of all higher-order G and D-formulas defined inductively by the following rules:

$$G := \top \mid A \mid G_1 \wedge G_2 \mid G_1 \vee G_2 \mid \exists x\, G,$$
$$D := A_r \mid G \supset A_r \mid D_1 \wedge D_2 \mid \forall x\, D.$$

Note that the type of quantified variables in this definition can be at *any* type including higher-order (predicate) types. The formulas of $\mathcal{D}_2$ are called *higher-order Horn clauses.* Notice that $\mathcal{G}_2$ is precisely the set of formulas contained in the set of terms $\mathcal{H}_1$.

Proving that uniform proofs are complete for the higher-order version of Horn clauses has some difficult parts. In particular, higher-order substitutions

into higher-order Horn clause may no longer be Horn clauses. Nadathur was able to prove [79, 82], however, that whenever there was a proof involving higher-order Horn clauses, it was also possible to restrain the higher-order (predicate) substitutions so that the only instances of Horn clauses were other Horn clauses.

Influence: Proof theory on logic programming

Sequent calculus provides a flexible framework for formalizing logic programming using Horn clauses with both first-order and higher-order quantification.

Notice that goal formulas (G-formulas in the definitions above) are not necessarily limited to atomic formulas: in the Horn clause setting, they can also be conjunctions, disjunctions, and existential quantifiers. Thus, backchaining is not the only inference rule that can be used in this setting. In fact, one can prove the following: when a sequent contains a non-atomic right-hand side (i.e., a goal formula with a logical connective) then the proof of that sequent can be assumed to be a right-introduction rule. Thus, provability with respect to this presentation of Horn clauses builds proofs divided into two phases: when the goal formula is atomic, the backchaining inference rule is used but when the goal formula is non-atomic, then the goal is reduced by using a right-introduction rule (reading proofs from the conclusion to premises).

This two-phase aspect of proof search has a natural appeal. The processing of logical connectives in the goal is fixed (by the right introduction rules). It is only when a non-logical symbol (the predicate at the head of an atomic formula) is encountered as the goal that we need to consult the (logic) program.

Now that we have a firm basis for logic programming using Horn clauses in sequent calculus, we can ask a natural question: What is the dynamics of proof search? More precisely, if $\mathcal{P} \vdash A$ is a root of a sequent calculus proof and $\mathcal{P}' \vdash A'$ is any other sequent in that proof (where A and A' are atomic formulas), then how are $\mathcal{P}$ and $\mathcal{P}'$, and A and A' related. In the case of Horn clauses, we know that there are rather natural proof systems for classical logic in which $\mathcal{P} = \mathcal{P}'$. Thus, during the search for a goal, there is no change to the left and, thus, the logic program is global and flat: every part of it is present at all times. Another way to describe this is to say that the only dynamics—the changing of atomic formulas—takes place within *non-logical* contexts, that is, in the scope of the non-logical symbols that are the predicates of atoms. Putting the dynamics of computation outside of logical contexts certainly seems to diminish the potential of logic to encode and reason about computational dynamics.

This characterization of Horn clauses has important implications for the structuring of programs: if a program clause is ever needed during a computation, it must be available at the beginning of that computation. Thus, Horn clauses do not support directly any hiding of one part of a program from

other parts of a program: such a lack is a significant problem for a modern programming language [62].

6 What's past is PROLOGue: intuitionistic logic extensions

Working from this last observation about how the left-hand-side of sequents using Horn clauses is a fixed and global value, the simple suggestion to use goals that are implications would allow contexts to grow as one moves up a proof from the conclusion to premises. In particular, Gentzen's right introduction rule for implication

$$\frac{\mathcal{P}, D \vdash G}{\mathcal{P} \vdash D \supset G}$$

can be interpreted as adding the new program element D (which might be a Horn clause) to the logic program $\mathcal{P}$. Thus, an attempt to prove the query $(D_1 \supset A_1) \wedge (D_2 \supset A_2)$ from the logic program $\mathcal{P}$ would be expected to yield the attempts to prove A_1 from $\mathcal{P} \cup \{D_1\}$ and to prove A_2 from $\mathcal{P} \cup \{D_2\}$. Thus, attempts to prove the two goals A_1 and A_2 are performed with different logic programs.

While an approach to adding a form of modularity to logic programming is rather immediate, one must confront the fact that classical logic does not provide the proper foundations for this notion of modularity. For example, one expects that attempting to prove $(D_1 \supset A_1) \vee (D_2 \supset A_2)$ from $\mathcal{P}$ would result in an attempt to prove A_1 from $\mathcal{P} \cup \{D_1\}$ or to prove A_2 from $\mathcal{P} \cup \{D_2\}$. But this interpretation is not supported by classical logic. Since the classical interpretation of the implication $D \supset G$ is the same as $(\neg D) \vee G$ then $(D_1 \supset A_1) \vee (D_2 \supset A_2)$ is logically equivalent to both $(D_2 \supset A_1) \vee (D_1 \supset A_2)$ and $(D_1 \supset (D_2 \supset (A_1 \vee A_2)))$. That is, classical logic does not support the intended scoping interpretation.

In the mid-1980's, the author was developing just such a scheme for providing λProlog [72, 81] with a form of modularity: the theory quickly settled on the need to use intuitionistic logic and not classical logic in order to achieve this approach to modularity [59, 61]. By the mid-1980's, intuitionistic logic and its proof theory had had a long development, much of that was in the general area of the Curry-Howard Isomorphism (proofs-as-programs). As it turns out, at about this same time, there was nearly simultaneous development of computational uses of large parts of intuitionistic logic that fell outside the Curry-Howard Isomorphism and more squarely in the proof-search framework. These various developments include the following.

- The N-Prolog language of Gabbay and Reyle [25, 26] was designed to allow hypothetical implications in a Prolog-like setting.
- McCarty [57, 58] explored using intuitionistic logic to extend the expressiveness of logic programs.
- Miller, Nadathur, Pfenning, and Scedrov [73, 74] developed a higher-order version of *hereditary Harrop formulas* in order to support rich notions

of abstractions—such as modules, abstract datatypes, and higher-order programming–within the logic programming setting.
- Paulson employed an intuitionistic logic to maintain proof states within the Isabelle theorem prover [87].
- Hällnais and Schroeder-Heister applied some proof-theoretic considerations to extend Horn clause programming in ways similar to these other approaches [32].

The simultaneous development of similar uses of intuitionistic logic within the logic programming (proof search) setting provided a great deal of confidence that intuitionistic logic and formulas with logical complexity much richer than Horn clauses could have important applications in computational logic. Since resolution refutations fundamentally rely on classical logic principles, the familiar framework on SLD-resolution needed to be rejected as a framework for these newly extended logic programming proposals. The sequent calculus provided just such a new starting point.

Influence: Proof theory on logic programming

The sequent calculus provided a simple and direct characterization of goal-directed proof search and to define a notion of *abstract logic programming languages*.

A *uniform proof* [73, 74] is a single conclusion (cut-free) sequent proof in which each occurrence of a sequent whose right-hand side contains a non-atomic formula is the conclusion of a right-introduction rule. In other words, a uniform proof is a sequent proof such that, for each occurrence of a sequent $\Gamma \vdash G$ in it, the following conditions are satisfied:

1. If G is $\top$ then that sequent is immediately proved.
2. If G is $B \wedge C$ then that sequent is inferred from $\Gamma \vdash B$ and $\Gamma \vdash C$.
3. If G is $B \vee C$ then that sequent is inferred from either $\Gamma \vdash B$ or $\Gamma \vdash C$.
4. If G is $\exists x\, B$ then that sequent is inferred from $\Gamma \vdash [t/x]B$ for some term t.
5. If G is $B \supset C$ then that sequent is inferred from $B, \Gamma \vdash C$.
6. If G is $\forall x\, B$ then that sequent is inferred from $\Gamma \vdash [c/x]B$, where c is an eigenvariable (parameter) that does not occur in the given sequent.

The notion of a uniform proof reflects the search instructions associated with the logical connectives. The logic program is only examined (via left-introduction rules) in the case that a non-logical symbol rises to the top of the query: such non-logical symbols are predicates and these are given meaning (axiomatized) by the logic program on the left-hand-side of a sequent. An *abstract logic programming language* is a triple $\langle \mathcal{D}, \mathcal{G}, \vdash \rangle$ such that for all finite subsets $\mathcal{P}$ of $\mathcal{D}$ and all formulas G of $\mathcal{G}$, $\mathcal{P} \vdash G$ holds if and only if there is a uniform proof of G from $\mathcal{P}$. It is in the following sense that uniform proofs are intended to capture the notion of *goal-directed search*. The impact on the

search for proofs is fixed by the top-level logical connective of the goal. We only examine the program when there is a non-logical symbol at the head of the sequent.

One example of an abstract logic programming languages is the one based on Horn clauses. In particular, both the triples $\langle \mathcal{D}_1, \mathcal{G}_1, \vdash_C \rangle$ (capturing first-order Horn clauses) and $\langle \mathcal{D}_2, \mathcal{G}_2, \vdash_C \rangle$ (capturing higher-order Horn clauses), are abstract logic programming languages. This statement is also true if $\vdash_C$ is replaced by $\vdash_I$ in both of these triples.

A more complex example is given by the following example. Let A be a syntactic variable that ranges over first-order atomic formulas. Let $\mathcal{G}_3$ and $\mathcal{D}_3$ be the sets of all first-order G- and D-formulas defined by the following rules:

$$G := \top \mid A \mid G_1 \wedge G_2 \mid G_1 \vee G_2 \mid \forall x\, G \mid \exists x\, G \mid D \supset G,$$
$$D := A \mid G \supset A \mid \forall x\, D \mid D_1 \wedge D_2.$$

Formulas in $\mathcal{D}_3$ are called *first-order hereditary Harrop* formulas. It is proved in [73] that the triple $\langle \mathcal{D}_3, \mathcal{G}_3, \vdash_I \rangle$ is an abstract logic programming language.

Let $\mathcal{H}_2$ be the set of all λ-normal terms that do not contain occurrences of the logical constants $\supset$ and $\bot$. Let A and A_r be syntactic variables denoting, respectively, atomic formulas and rigid atomic formulas in $\mathcal{H}_2$. Let $\mathcal{G}_4$ and $\mathcal{D}_4$ be the sets of G- and D-formulas that are defined by the following mutual recursion:

$$G := \top \mid A \mid G_1 \wedge G_2 \mid G_1 \vee G_2 \mid \forall x\, G \mid \exists x\, G \mid D \supset G$$
$$D := A_r \mid G \supset A_r \mid \forall x\, D \mid D_1 \wedge D_2.$$

The formulas of $\mathcal{D}_4$ are called *higher-order hereditary Harrop* formulas and it is proved in [73] that the triple $\langle \mathcal{D}_4, \mathcal{G}_4, \vdash_I \rangle$ is an abstract logic programming language.

As in the case of Horn clauses, proof search with hereditary Harrop formulas yields uniform proofs that are organized into alternating phases: one phase reduces goal formulas (using right-introduction rules), and one phase performs backchaining steps (using left-introduction rules and the initial rule) [63, 73].

The λProlog programming language was designed to implement most of the intuitionistic theory of higher-order hereditary Harrop formulas: a key design goal of that language was to test the abstraction mechanisms that that logic provides [72, 81]. Since there is a significant gap between having a description of a logic programming language in a sequent calculus and an actual implementation of that language, there were a number of significant developments that needed to be made prior to having comprehensive implementations of that language, of which there are two currently, namely Teyjus [83, 89] and ELPI [20]. The description of a unification algorithm that works well in the sequent calculus where eigenvariables are present was one of those challenges [63, 64, 80]: such unification made it possible to avoid the use of Skolem terms.

At the end of Section 5 we described the dynamics of proof search with Horn clauses as *flat* since the logic program used during proof search never changes during a computation. When we examine the dynamics of change

using hereditary Harrop formulas, we note that the left-hand side of sequents (the logic program) can grow monotonically as we move from the conclusion to premises.

The overview of structuring mechanisms for logic programming given in [11] provides still additional examples of how proof theory considerations can provide or can influence this aspect of designing logic programming languages.

7 Linear Logic and Logic Programming

As we noted in the previous section, the use of intuitionistic logic and hereditary Harrop formulas allowed logic programs to be seen as a structure that grows in a stack-based discipline as the search for proofs moves from conclusion to premises. While such growth in logic programs is an improvement over what was available using only Horn clauses, many additional problems existed in computational logic that were just-out-of-reach of having an elegant solution using intuitionistic logic.

For example, in the area of natural language, a good treatment of *filler-gap dependencies* (used to characterize such crucial natural language constructs as questions and relative clauses) was hard to achieve using standard Horn clause-based logic grammars and lead to the development of the slashed non-terminal in the framework of Generalized Phrase Structure Grammar (GPSG) [27]. A different approach using intuitionistic logic made it possible to identify the linguistic notion of *gap introduction* as hypothesis introduction using an implicational goal. As reported in by Pareschi [85, 86], that technique provided an elegant new perspective to that linguistic phenomenon but it also failed to treat known restrictions on the distribution and use of gaps-as-hypotheses: in particular, gaps needed to be used and they could not appear in certain parts of phrases.

For another example, Hodas [38, 40] described how it was possible to capture partially the notion of objects-with-state within logic programming. Again, intuitionistic logic provides a partial solution. In particular, it is possible to store the value of a register as an atomic formula among the other clauses of a logic program. For example, the atomic formula $reg(4)$ can encode the fact that a register currently has value 4. Unfortunately, there is no way to have that atomic formula replaced with, say, $reg(5)$ within intuitionistic logic. More specifically, it is not possible to write a logic program clause such that backchaining on it would give rise to the following synthetic rule.

$$\frac{\mathcal{P}, reg(5) \vdash A'}{\mathcal{P}, reg(4) \vdash A}$$

where both A and A' are atomic formulas. The best one can do within intuitionistic logic is to move to a context in which both atoms $reg(4)$ and $reg(5)$ are present: that is, the following inference rule is possible.

$$\frac{\mathcal{P}, reg(4), reg(5) \vdash A'}{\mathcal{P}, reg(4) \vdash A}$$

Unfortunately, this situation (where a register has two different values) does not provide a proper modeling of a register.

With the appearance of Girard's linear logic [29], it was possible to extend the design of previous logic programming languages so that they could solve the cited problems in both the gap-threading and the state-encapsulation situations. Logic programming provided other important examples that helped to convince a number of computer scientists of the value of linear logic to computational logic: beyond the two examples mentioned above, additional examples appeared in the areas of concurrency [47, 65], Petri nets [21, 44], and theorem proving [39].

Influence: Logic programming on proof theory

A large set of examples arose from the logic programming community, particularly those working with hypothetical queries, in which linear logic was immediately applicable. Examples such as these increased the confidence in the utility of linear logic in computational logic.

Linear logic also provided a richer analysis of the role of structural rules in Gentzen's sequent calculus and, as a result, greatly improved our understanding of proof search. For example, if one can restrict the uses of the structural rule of contraction (which can be done in linear logic), one can often turn a naive proof search mechanism into a complete decision procedure.

After Girard's introduction of linear logic in 1987, it became clear that there should be *linear logic programming languages*: the logic programming paradigm lacked certain features (e.g., side-effects and communications) which linear logic seems capable of capturing.

Among the first linear logic programming languages designed, there was a divergence along two axes. One of the most challenging connectives in linear logic for computer scientists to appreciate was the *multiplicative disjunction* $\parr$. For the proof theorist, this connective was not a challenge since it could be identified with the comma appearing on the right of Gentzen's multiple conclusion sequents. It could also be seen as the de Morgan dual of the multiplicative conjunction $\otimes$. In computational logic, however, intuitions coming from intuitionistic logic can make it difficult to find computational meaning for $\parr$ since Gentzen identified intuitionistic logic with single conclusion sequents. One early development thread in linear logic programming languages also avoided using $\parr$. However, the first linear logic programming language made prominent use of that connective.

7.1 Linear Objects

Historically speaking, the first proposal for a linear logic programming language was LO (Linear Objects) by Andreoli and Pareschi [4, 5]. LO is an extension to the Horn clause paradigm in which, roughly speaking, the role

$$\frac{}{\Psi;\Delta \longrightarrow \top} \qquad \frac{\Psi;\Delta \longrightarrow G_1 \qquad \Psi;\Delta \longrightarrow G_2}{\Psi;\Delta \longrightarrow G_1 \mathbin{\&} G_2}$$

$$\frac{\Psi, G_1;\Delta \longrightarrow G_2}{\Psi;\Delta \longrightarrow G_1 \Rightarrow G_2} \qquad \frac{\Psi;\Delta, G_1 \longrightarrow G_2}{\Psi;\Delta \longrightarrow G_1 \multimap G_2} \qquad \frac{\Psi;\Delta \longrightarrow B[y/x]}{\Psi;\Delta \longrightarrow \forall_\tau x.B}\ \dagger$$

$$\frac{\Psi, D;\Delta \xrightarrow{D} A}{\Psi, D;\Delta \longrightarrow A} \qquad \frac{\Psi;\Delta \xrightarrow{D} A}{\Psi;\Delta, D \longrightarrow A} \qquad \frac{}{\Psi;\cdot \xrightarrow{A} A}$$

$$\frac{\Psi;\Delta \xrightarrow{D_i} A}{\Psi;\Delta \xrightarrow{D_1 \wedge D_2} A} \qquad \frac{\Psi;\cdot \longrightarrow G \qquad \Psi;\Delta \xrightarrow{D} A}{\Psi;\Delta \xrightarrow{G \Rightarrow D} A}$$

$$\frac{\Psi;\Delta_1 \longrightarrow G \qquad \Psi;\Delta_2 \xrightarrow{D} A}{\Psi;\Delta_1,\Delta_2 \xrightarrow{G \multimap D} A} \qquad \frac{\Psi;\Delta \xrightarrow{D[t/x]} A}{\Psi;\Delta \xrightarrow{\forall_\tau x.D} A}\ \ddagger$$

Fig. 2 The proof system for Lolli. The rule for universal quantification has the proviso † that y is not free in any formula of the conclusion. In the $\forall$-left rule, the proviso ‡ requires t to be a term of type τ. The completeness of this proof system for Lolli is given in [39].

of atomic formulas in Horn clauses is generalized to multisets (built using $⅋$) of atomic formulas. In LO, backchaining captures multiset rewriting and the dominant examples of LO were taken from those domains where multiset rewriting had proved useful, namely, object-oriented programming and the coordination of processes. Program clauses in LO are formulas of the form

$$\forall \bar{y}(G_1 \hookrightarrow \cdots \hookrightarrow G_m \hookrightarrow (A_1 ⅋ \cdots ⅋ A_p)).$$

Here $p > 0$ and $m \geq 0$; occurrences of $\hookrightarrow$ are either occurrences of $\multimap$ (linear implication) or $\Rightarrow$ (intuitionistic implication); $G_1, \ldots G_m$ are built from $\bot$, $⅋$, $?$, $\top$, $\&$, and $\forall$; and $A_1, \ldots A_m$ are atomic formulas. The two implications are related by the familiar linear logic equivalence between $B \Rightarrow C$ and $(!\,B) \multimap C$. By applying uncurrying, the displayed formula above can be written as

$$\forall \bar{y}[(\S G_1 \otimes \cdots \otimes \S G_m) \multimap (A_1 ⅋ \cdots ⅋ A_p)]$$

where $\S G_i$ is either G_i if G_i is to the immediate left of a $\multimap$ or is $!\,G_i$ if G_i is to the immediate left of a $\Rightarrow$. Note that if this displayed formula contained no occurrences of $⅋$ and $\Rightarrow$ then it is an easy matter to view that formula as a simple Horn clause.

7.2 Lolli

The Lolli logic programming language was introduced by Hodas and the author as a linear logic extension to the intuitionistic theory of hereditary Harrop formulas. In particular, Lolli can be seen as a revision and small extension to the logic of hereditary Harrop formulas (Section 6). The description of goal formulas and program clauses can be given as follows.

$$G := A \mid G_1 \mathbin{\&} G_2 \mid G_1 \otimes G_2 \mid G_1 \oplus G_2 \mid \forall x\, G \mid \exists x\, G \mid D \Rightarrow G \mid D \multimap G,$$
$$D := A \mid G \multimap A \mid G \Rightarrow A \mid \forall x\, D \mid D_1 \mathbin{\&} D_2.$$

The difference between these formulas and those associated with the hereditary Harrop formulas is that the conjunction in intuitionistic logic could be mapped to either $\otimes$ or &: for goal formulas, both options are available while for program clauses, only the & option is available. (Technically, the conjunctions and disjunction should also be accompanied by their units.)

A proof system for Lolli is presented in Figure 2. This sequent calculus proof system differs from those used by Gentzen in [28] and Girard in [29] in that (i) the left-hand context is divided into two parts $\Psi;\Delta$ (where both Ψ and Δ are multisets of D formulas) and (ii) there are two kinds of sequents written as $\Psi;\Delta \longrightarrow G$ and $\Psi;\Delta \xrightarrow{D} A$. The context Ψ denotes those formulas that can be used any number of times during the search for a proof while those in Δ are controlled using the restrictions on the structural rules of contraction and weakening. As a result, the context Ψ is often called the "unbounded" context and Δ is often called the "bounded" or "linear" context. These sequents can be mapped into the more usual linear logic sequents by rewriting $\Psi;\Delta \longrightarrow G$ as $!\Psi,\Delta \vdash G$ and $\Psi;\Delta \xrightarrow{D} A$ as $!\Psi,\Delta,D \vdash A$. The formula that is placed on top of the turnstile in some of these sequents is the formula involved with backchaining: note that the left-introduction rules are only applied to the formula that labels such a turnstile.

We can now illustrate how we can model the change in a register's value. Assume that Ψ contains the formula

$$D = \forall N \forall G [\mathit{reg}(N) \multimap (\mathit{reg}(N+1) \multimap G) \multimap \mathit{inc}(G)]$$

Using the proof rules in Figure 2 we can write the following partial derivation.

$$\dfrac{\dfrac{\dfrac{}{\Psi;\cdot \xrightarrow{\mathit{reg}(4)} \mathit{reg}(4)}}{\Psi;\mathit{reg}(4) \longrightarrow \mathit{reg}(4)} \qquad \dfrac{\dfrac{\Psi;\Delta,\mathit{reg}(5) \longrightarrow G}{\Psi;\Delta \longrightarrow \mathit{reg}(5) \multimap G} \quad \dfrac{}{\Psi;\cdot \xrightarrow{\mathit{inc}(G)} \mathit{inc}(G)}}{\Psi;\Delta \xrightarrow{D'} \mathit{inc}(G)}}{\dfrac{\Psi;\mathit{reg}(4) \xrightarrow{D} \mathit{inc}(G)}{\Psi;\Delta,\mathit{reg}(4) \longrightarrow \mathit{inc}(G)}}$$

Here, D' is the formula $(\mathit{reg}(5) \multimap G) \multimap \mathit{inc}(G)$. Thus, the synthetic inference rule that this derivation gives rise to is simply

$$\dfrac{\Psi;\Delta,\mathit{reg}(5) \longrightarrow G}{\Psi;\Delta,\mathit{reg}(4) \longrightarrow \mathit{inc}(G)}$$

Critical for the correct modeling of the change in state of this register is the splitting of the linear context in the $\multimap$ left introduction rule between its two premises and the fact that the linear context must be empty in the initial rules.

7.3 Goal-directed search with multiple conclusion

A natural question is whether or not it is possible to view LO and the Lolli as sublanguages of a larger linear logic programming language. While Lolli contains occurrences of many linear logic connectives, it does not allow occurrences of $\invamp$, its unit $\bot$, and its associated exponential ?. One thing to note is that if one adds to Lolli just $\bot$, *all* connectives of linear logic can then be defined. For example, $B \invamp C$ can be defined as $(B \multimap \bot) \multimap C$ and $?B$ can be defined as $(B \multimap \bot) \Rightarrow \bot$. In [66, 67], the author proposed a new logic programming language, called Forum, which results from adding $\bot$, $\invamp$, and ? to Lolli. Thus, Forum is essentially a presentation of all of linear logic as a logic programming language.

The most direct way to view all of linear logic as a logic programming language suggests attempting to generalize the notion of uniform proof from single-conclusion to multiple-conclusion sequents. This can be done if we insist that goal-reduction should continue to be independent of not only the logic program but also other goals, *i.e.*, multiple goals should be reducible *simultaneously.* Although the sequent calculus does not directly allow for simultaneous rule application, it can be simulated easily by referring to permutations of inference rules [46]. In particular, we can require that if two or more right-introduction rules can be used to derive a given sequent, then all possible orders of applying those right-introduction rules can be obtained from any other order simply by permuting right-introduction inferences. It is easy to see that the following definition of uniform proofs for multiple-conclusion sequents generalizes that for single-conclusion sequents: a cut-free, sequent proof Ξ is *uniform* if for every subproof Ψ of Ξ and for every non-atomic formula occurrence B in the right-hand side of the end-sequent of Ψ, there is a proof Ψ' that is equal to Ψ up to permutation of inference rules and is such that the last inference rule in Ψ' introduces the top-level logical connective occurring in B [65, 67]. The notion of an abstract logic programming language can be generalized to include this extended notion of uniform proof.

7.4 Focusing

As it turns out, the completeness of multiple conclusion uniform proofs for Forum had actually been solved a couple years before the introduction of Forum. The Ph.D. dissertation of Andreoli [2] introduced a new sequent calculus proof system for linear logic, called a *focused proof system*, that was composed of two kinds of sequents and two phases of proof construction. That proof system resembles the proof system in Figure 2 and the formula that is placed over the turnstile in that figure corresponds to the focus that exists in one of the phases of the focused proof. The completeness of focused proofs (see also [3]) provided the completeness result for Forum [67] (see also [10]). However, Andreoli's presentation of a focused proof system of linear logic provided important and deeper insights into the structure of proof search in the sequent

calculus. In particular, Andreoli's analysis of the two phases of rule application was based on a notion of *polarity* of logical connectives and that polarity was flipped by de Morgan duality. The use of two phases of proof construction was a powerful addition to the results of pure proof theory. Several subsequent efforts have been made to provide focused proof systems for classical and intuitionistic logic all of which appear to be captured by the LKF and LJF focused proofs system of Liang and the author [53].

Influence: Logic programming on proof theory

When the notion of uniform proof, with its two phases for structuring proof search, were extended to linear logic, a far richer analysis of proof structure was developed using *focused proofs*.

7.5 Other linear logic programming languages

Besides LO, Lolli, and Forum, various other subsets of linear logic have been studied as logic programming languages. The Lygon system of Harland and Pym [35] was based on a notion of multiple-conclusion goal-directed proof search different from the one described above [88]. The operational semantics for proof search in Lygon is different and more complex than the alternating of goal-reduction and backchaining found in, say, Forum. Various other specification logics have also been developed, often designed directly to deal with particular application areas. In particular, the language ACL by Kobayashi and Yonezawa [47, 48] captures simple notions of asynchronous communication by identifying the send and read primitives with two complementary linear logic connectives. Lincoln and Saraswat have developed a linear logic version of concurrent constraint programming [54, 93] and Fages, Ruet, and Soliman have analyzed similar extensions to the concurrent constraint paradigm [22, 91].

Let G and H be formulas composed of $\bot$, $\invamp$, and $\forall$. Closed formulas of the form $\forall\bar{x}[G \multimap H]$ (where H is not $\bot$) have been called *process clauses* in [65] and are used there to encode a calculus similar to the π-calculus: the universal quantifier in goals are used to encode name restriction. These clauses, when written in their contrapositive form (replacing, for example, $\invamp$ with $\otimes$), have been called *linear Horn clauses* by Kanovich and have been used to model computation via multiset rewriting [43]. A generalization of process clauses was presented in [68] and was applied to the description of security protocols.

Some aspects of dependent typed λ-calculi overlap with notions of abstract logic programming languages. Within the setting of intuitionistic, single-side sequents, uniform proofs are similar to $\beta\eta$-long normal forms in natural deduction and typed λ-calculus. The LF logical framework [36] can be mapped naturally [23] into a higher-order extension of hereditary Harrop formulas [73]. Inspired by such a connection and by the design of Lolli, Cervesato and Pfenning developed a linear extension to LF called Linear LF [13, 14].

An overview of research into the design of linear logic programming language up until 2004 can be found in [70].

Influence: Proof theory on logic programming

Linear logic allowed for the rich, new logic programming languages that modularly extend previously designed logic programming languages.

8 First-order and higher-order quantification

While most work in proof theory and logic programming has addressed only first-order quantification, several researchers have defined and implemented logic programming languages that include higher-order quantification.

Church, the inventor of the λ-calculus, is the inventor of the most popular version of higher-order logic in use in computational logic presently. In particular, Church's Simple Theory of Types (STT) [16] defines the syntax of both terms and formulas using simply typed λ-terms (simple types have been introduced in Section 3). STT used only one form of binding and that is the one used to form λ-abstractions: all other bindings—for example, the universal and existential quantifier—are built using the λ-binder. In STT, it was possible to quantify over variables of primitive type (first-order quantification) as well as types containing the arrow constructor and the primitive type o, such as $\mathit{list} \to \mathit{list}$, $\mathit{nat} \to o$, and $(\mathit{list} \to o) \to o$ (higher-order quantification).

When implementing computer systems that need to manipulate syntactic expressions in artificial and natural languages, the strings containing those syntactic expressions need to be parsed. The result of such a parse is generally a *parse tree* or *abstract syntax tree* representation capturing the structure of the parsed expression. Most traditional programming language—functional, imperative, logic—have convenient and flexible means to process tree structures. However, a majority of syntactic expressions that need to be parsed and manipulated contain more than recursive tree structures: they also contain binding structures. While binding structures can, of course, be encoded in tree structures (using techniques such as de Bruijn's nameless dummies [9]) no traditional programming language contains direct support for such an important feature of most syntactic expressions.

A good option for properly merging Church's logic with Gentzen's sequent calculus is to continue Church's identification of bindings to one additional level. That is, term-level bindings (λ-abstractions) and formula-level bindings (quantifiers) need to also be merged with proof-level bindings, which are the eigenvariables of the sequent calculus. It is possible to consider eigenvariables to be bindings around sequents: that is, if $\mathcal{V}$ is a set of distinct variables then the expression $\mathcal{V}\colon \Gamma \vdash \Delta$ can be interpreted as the formal binding of the variables in $\mathcal{V}$ over the formulas in both Γ and Δ.

To illustrate this kind of merging of bindings at these three different levels, consider specifying the binary predicate *typeof* whose arguments are encodings

of an untyped λ-term and an encoding of a simple type. The intended meaning of this predicate is that (*typeof* $\lceil B\rceil$ $\lceil\tau\rceil$) holds if and only if the untyped λ-term B can be typed with τ. For this example, we will write $\lceil t\rceil$ to denote some encoding of untyped λ-terms into simply typed terms: the key for this encoding is that bindings in the untyped terms are encoded as binders in the encoded terms. We also assume that there is some encoding, also written $\lceil\tau\rceil$, of simple type expressions into (first-order) terms. Part of the specification of *typeof* is given by the following derivation.

$$\dfrac{\dfrac{\mathcal{V}, x : \Delta, \mathit{typeof}\ x\ \lceil\alpha\rceil \vdash \mathit{typeof}\ \lceil B\rceil\ \lceil\beta\rceil}{\mathcal{V} : \Delta \vdash \forall x(\mathit{typeof}\ x\ \lceil\alpha\rceil \supset \mathit{typeof}\ \lceil B\rceil\ \lceil\beta\rceil)}\ \forall R}{\mathcal{V} : \Delta \vdash \mathit{typeof}\ \lceil\lambda x.B\rceil\ \lceil\alpha \to \beta\rceil}$$

In this case, the binder named x moves from *term-level* (λx) to *formula-level* ($\forall x$) to *proof-level* (as an eigenvariable in $\mathcal{V}, x$). Thus an integration of Church's STT and Gentzen's sequent calculus provides for the *mobility of binders*.

Influence: Logic programming on proof theory

Specifications in logic programs that deal with bindings within data required that bindings move from terms, to formulas, to proofs. Proof theory needed to account for this *mobility* of binders by identifying eigenvariables as proof-level binders.

As it turns out, λProlog was the first programming language with such support since it embraced this notion of *binder mobility* [69, 71]. Such a notion exists in various other specification languages. For example, dependently typed λ-calculus makes similar identifications, although often in the setting of natural deduction [19, 36]. It is also a feature of the Isabelle theorem prover in which the technical device of $\forall$-lifting is used to link eigenvariables to formula-level quantification.

9 Conclusion

There have been significant reciprocal influences between researchers working on structural proof theory and those working on logic programming. While it is not surprising to find that the older and more mature topic of proof theory provided the bulk of that influence, it is still the case that problems identify within logic programming forced proof theorists to deepen and extend their results. The development of two-phase proof constructions that resulted in focused proof systems might be the most prominent example; the encoding of binder mobility has been a second such example.

Acknowledgements I thank the participants of the "Fourth Symposium on the History and Philosophy of Programming" (HaPoP 2018) meeting and the ANR project PROGRAMme (ANR-17-CE38-0003-01) for their comments on a talk based on an earlier version of this paper.

References

1. Aït-Kaci, H.: Warren's Abstract Machine: A Tutorial Reconstruction. Logic Programming Research Reports and Notes. MIT Press, Cambridge, MA (1991). URL http://wambook.sourceforge.net/
2. Andreoli, J.M.: Proposal for a synthesis of logic and object-oriented programming paradigms. Ph.D. thesis, University of Paris VI (1990)
3. Andreoli, J.M.: Logic programming with focusing proofs in linear logic. J. of Logic and Computation **2**(3), 297–347 (1992). DOI 10.1093/logcom/2.3.297
4. Andreoli, J.M., Pareschi, R.: Communication as fair distribution of knowledge. In: Proceedings of OOPSLA 91, pp. 212–229 (1991)
5. Andreoli, J.M., Pareschi, R.: Linear objects: Logical processes with built-in inheritance. New Generation Computing **9**(3-4), 445–473 (1991)
6. Andrews, P.B.: Resolution in type theory. J. of Symbolic Logic **36**, 414–432 (1971)
7. Apt, K.R., van Emden, M.H.: Contributions to the theory of logic programming. J. of the ACM **29**(3), 841–862 (1982)
8. Barendregt, H.: The impact of the lambda calculus in logic and computer science. Bulletin of Symbolic Logic **3**(2), 181–215 (1997)
9. de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with an application to the Church-Rosser theorem. Indagationes Mathematicae **34**(5), 381–392 (1972)
10. Bruscoli, P., Guglielmi, A.: On structuring proof search for first order linear logic. Theoretical Computer Science **360**(1-3), 42–76 (2006)
11. Bugliesi, M., Lamma, E., Mello, P.: Modularity in logic programming. Journal of Logic Programming **19/20**, 443–502 (1994)
12. Burns, A.R.: Programming in OCCAM 2. Addison-Wesley (1988)
13. Cervesato, I., Pfenning, F.: A linear logic framework. In: 11th Symp. on Logic in Computer Science, pp. 264–275. IEEE Computer Society Press, New Brunswick, New Jersey (1996)
14. Cervesato, I., Pfenning, F.: A Linear Logical Framework. Information & Computation **179**(1), 19–75 (2002)
15. Chen, W., Kifer, M., Warren, D.S.: HILOG: a foundation for higher-order logic programming. J. of Logic Programming **15**(3), 187–230 (1993)
16. Church, A.: A formulation of the Simple Theory of Types. J. of Symbolic Logic **5**, 56–68 (1940). DOI 10.2307/2266170
17. Church, A.: The Calculi of Lambda Conversion. (AM-6) (Annals of Mathematics Studies). Princeton University Press, Princeton, NJ, USA (1985)
18. Colmerauer, A., Roussel, P.: The birth of Prolog. SIGPLAN Notices **28**(3), 37–52 (1993)
19. Coquand, T., Huet, G.: The calculus of constructions. Information and Computation **76**(2/3), 95–120 (1988)
20. Dunchev, C., Guidi, F., Coen, C.S., Tassi, E.: ELPI: fast, embeddable, λProlog interpreter. In: M. Davis, A. Fehnker, A. McIver, A. Voronkov (eds.) Logic for Programming, Artificial Intelligence, and Reasoning -

20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings, *LNCS*, vol. 9450, pp. 460–468. Springer (2015). DOI 10.1007/978-3-662-48899-7_32. URL http://dx.doi.org/10.1007/978-3-662-48899-7_32
21. Engberg, U., Winskel, G.: Petri nets and models of linear logic. In: A. Arnold (ed.) CAAP'90, LNCS 431, pp. 147–161. Springer (1990)
22. Fages, F., Ruet, P., Soliman, S.: Phase semantics and verification of concurrent constraint programs. In: V. Pratt (ed.) 13th Symp. on Logic in Computer Science. IEEE (1998)
23. Felty, A.: Transforming specifications in a dependent-type lambda calculus to specifications in an intuitionistic logic. In: G. Huet, G.D. Plotkin (eds.) Logical Frameworks. Cambridge University Press (1991)
24. Fitting, M.: Resolution for intuitionistic logic. In: Proc. International Symposium of Methodologies for Intelligent Systems ISMIS'87, pp. 400–407 (1987)
25. Gabbay, D.M.: N-Prolog: An extension of Prolog with hypothetical implication II—logical foundations, and negation as failure. Journal of Logic Programming **2**(4), 251–283 (1985)
26. Gabbay, D.M., Reyle, U.: N-Prolog: An extension of Prolog with hypothetical implications. I. Journal of Logic Programming **1**, 319–355 (1984)
27. Gazdar, G., Klein, E., Pullum, G., Sag, I.: Generalized Phrase Structure Grammar. Harvard University Press, Cambridge, Massachusetts (1985)
28. Gentzen, G.: Investigations into logical deduction. In: M.E. Szabo (ed.) The Collected Papers of Gerhard Gentzen, pp. 68–131. North-Holland, Amsterdam (1935). DOI 10.1007/BF01201353
29. Girard, J.Y.: Linear logic. Theoretical Computer Science **50**(1), 1–102 (1987). DOI 10.1016/0304-3975(87)90045-4
30. Gödel, K.: On formally undecidable propositions of the principia mathematica and related systems. I. In: M. Davis (ed.) The Undecidable. Raven Press (1965)
31. Green, C.: Theorem proving by resolution as a basis for question-answering systems. Machine Intelligence **4** (1969)
32. Hallnäs, L., Schroeder-Heister, P.: A proof-theoretic approach to logic programming. II. Programs as definitions. J. of Logic and Computation **1**(5), 635–660 (1991)
33. Halpern, J.Y., Harper, R., Immerman, N., Kolaitis, P.G., Vardi, M.Y., Vianu, V.: On the unusual effectiveness of logic in computer science. Bulletin of Symbolic Logic **7**(1), 213–236 (2001)
34. Hardy, G.H.: A mathematician's apology. Cambridge University Press (1940)
35. Harland, J., Pym, D., Winikoff, M.: Programming in Lygon: An overview. In: Proceedings of the Fifth International Conference on Algebraic Methodology and Software Technology, pp. 391–405 (1996)
36. Harper, R., Honsell, F., Plotkin, G.: A framework for defining logics. Journal of the ACM **40**(1), 143–184 (1993)
37. Hoare, C.A.R.: Communicating Sequential Processes. Prentice-Hall (1985)

38. Hodas, J., Miller, D.: Representing objects in a logic programming language with scoping constructs. In: D.H.D. Warren, P. Szeredi (eds.) 1990 International Conference in Logic Programming, pp. 511–526. MIT Press (1990)
39. Hodas, J., Miller, D.: Logic programming in a fragment of intuitionistic linear logic. Information and Computation **110**(2), 327–365 (1994)
40. Hodas, J.S.: Logic programming in intuitionistic linear logic: Theory, design, and implementation. Ph.D. thesis, University of Pennsylvania, Department of Computer and Information Science (1994)
41. Huet, G.P.: A mechanization of type theory. In: Proceedings of the 3rd International Joint Conference on Artificial Intelligence, pp. 139–146 (1973)
42. Jaffar, J., Lassez, J.L.: Constraint logic programming. In: Proceedings of the 14th ACM Symposium on the Principles of Programming Languages (1987)
43. Kanovich, M.: The complexity of Horn fragments of linear logic. Annals of Pure and Applied Logic **69**, 195–241 (1994)
44. Kanovich, M.I.: Petri nets, Horn programs, Linear Logic and vector games. Annals of Pure and Applied Logic **75**(1–2), 107–135 (1995). DOI 10.1017/S0960129500001328
45. Kleene, S.C.: A theory of positive integers in formal logic, part I. American Journal of Mathematics **57**, 153–173 (1935)
46. Kleene, S.C.: Permutabilities of inferences in Gentzen's calculi LK and LJ. Memoirs of the American Mathematical Society **10**, 1–26 (1952)
47. Kobayashi, N., Yonezawa, A.: ACL - A concurrent linear logic programming paradigm. In: D. Miller (ed.) Logic Programming - Proceedings of the 1993 International Symposium, pp. 279–294. MIT Press (1993)
48. Kobayashi, N., Yonezawa, A.: Asynchronous communication model based on linear logic. Formal Aspects of Computing **3**, 279–294 (1994)
49. Kowalski, R.A.: Predicate logic as a programming language. Information Processing **74**, 569–574 (1974)
50. Kowalski, R.A., Kuehner, D.: Linear resolution with selection function. Artificial Intelligence **2**, 227–260 (1971)
51. Lambek, J., Scott, P.J.: Introduction to Higher Order Categorical Logic. Cambridge University Press (1986)
52. Lawvere, F.W.: Functorial semantics of algebraic theories. Proceedings National Academy of Sciences USA **50**, 869–872 (1963)
53. Liang, C., Miller, D.: Focusing and polarization in linear, intuitionistic, and classical logics. Theoretical Computer Science **410**(46), 4747–4768 (2009). DOI 10.1016/j.tcs.2009.07.041
54. Lincoln, P., Saraswat, V.: Higher-order, linear, concurrent constraint programming (1993). Ftp://parcftp.xerox.com/pub/ccp/lcc/hlcc.dvi
55. Lobo, J., Minker, J., Rajasekar, A.: Foundations of disjunctive logic programming. MIT Press (1992)
56. Martin-Löf, P.: Constructive mathematics and computer programming. In: C.A.R. Hoare, J.C. Shepherdson (eds.) Mathematical Logic and Programming Languages, pp. 167–184. Prentice-Hall (1985)

57. McCarty, L.T.: Clausal intuitionistic logic I. fixed point semantics. Journal of Logic Programming **5**, 1–31 (1988)
58. McCarty, L.T.: Clausal intuitionistic logic II. tableau proof procedure. Journal of Logic Programming **5**, 93–132 (1988)
59. Miller, D.: A theory of modules for logic programming. In: R.M. Keller (ed.) Third Annual IEEE Symposium on Logic Programming, pp. 106–114. Salt Lake City, Utah (1986)
60. Miller, D.: A compact representation of proofs. Studia Logica **46**(4), 347–370 (1987)
61. Miller, D.: A logical analysis of modules in logic programming. Journal of Logic Programming **6**(1-2), 79–108 (1989)
62. Miller, D.: Abstractions in logic programming. In: P. Odifreddi (ed.) Logic and Computer Science, pp. 329–359. Academic Press (1990). URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/AbsInLP.pdf.pdf
63. Miller, D.: A logic programming language with lambda-abstraction, function variables, and simple unification. J. of Logic and Computation **1**(4), 497–536 (1991)
64. Miller, D.: Unification under a mixed prefix. Journal of Symbolic Computation **14**(4), 321–358 (1992)
65. Miller, D.: The π-calculus as a theory in linear logic: Preliminary results. In: E. Lamma, P. Mello (eds.) 3rd Workshop on Extensions to Logic Programming, no. 660 in LNCS, pp. 242–265. Springer, Bologna, Italy (1993). URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/pic.pdf
66. Miller, D.: A multiple-conclusion meta-logic. In: S. Abramsky (ed.) 9th Symp. on Logic in Computer Science, pp. 272–281. IEEE Computer Society Press, Paris (1994)
67. Miller, D.: Forum: A multiple-conclusion specification logic. Theoretical Computer Science **165**(1), 201–232 (1996)
68. Miller, D.: Encryption as an abstract data-type: An extended abstract. In: I. Cervesato (ed.) Proceedings of FCS'03: Foundations of Computer Security, pp. 3–14 (2003)
69. Miller, D.: Bindings, mobility of bindings, and the ∇-quantifier. In: J. Marcinkowski, A. Tarlecki (eds.) 18th International Conference on Computer Science Logic (CSL) 2004, *LNCS*, vol. 3210, p. 24 (2004)
70. Miller, D.: Overview of linear logic programming. In: T. Ehrhard, J.Y. Girard, P. Ruet, P. Scott (eds.) Linear Logic in Computer Science, *London Mathematical Society Lecture Note*, vol. 316, pp. 119–150. Cambridge University Press (2004)
71. Miller, D.: Mechanized metatheory revisited. Journal of Automated Reasoning (2018). DOI 10.1007/s10817-018-9483-3. URL https://doi.org/10.1007/s10817-018-9483-3
72. Miller, D., Nadathur, G.: Programming with Higher-Order Logic. Cambridge University Press (2012). DOI 10.1017/CBO9781139021326
73. Miller, D., Nadathur, G., Pfenning, F., Scedrov, A.: Uniform proofs as a foundation for logic programming. Annals of Pure and Applied Logic **51**(1–2), 125–157 (1991)

74. Miller, D., Nadathur, G., Scedrov, A.: Hereditary Harrop formulas and uniform proof systems. In: D. Gries (ed.) 2nd Symp. on Logic in Computer Science, pp. 98–105. Ithaca, NY (1987)
75. Milner, R.: A Calculus of Communicating Systems, *LNCS*, vol. 92. Springer, New York, NY (1980)
76. Milner, R.: Communicating and Mobile Systems: The π-Calculus. Cambridge University Press, New York, NY, USA (1999)
77. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes, Part I. Information and Computation **100**(1), 1–40 (1992)
78. Minker, J., Seipel, D.: Disjunctive logic programming: A survey and assessment. In: Computational Logic: Logic Programming and Beyond, vol. 2407, pp. 472–511. Springer (2002)
79. Nadathur, G.: A higher-order logic as the basis for logic programming. Ph.D. thesis, University of Pennsylvania (1987)
80. Nadathur, G., Jayaraman, B., Kwon, K.: Scoping constructs in logic programming: Implementation problems and their solution. Journal of Logic Programming **25**(2), 119–161 (1995). DOI 10.1016/0743-1066(95)00037-K. URL https://doi.org/10.1016/0743-1066(95)00037-K
81. Nadathur, G., Miller, D.: An Overview of λProlog. In: Fifth International Logic Programming Conference, pp. 810–827. MIT Press, Seattle (1988).
82. Nadathur, G., Miller, D.: Higher-order Horn clauses. Journal of the ACM **37**(4), 777–814 (1990)
83. Nadathur, G., Mitchell, D.J.: System description: Teyjus — A compiler and abstract machine based implementation of λProlog. In: H. Ganzinger (ed.) 16th Conf. on Automated Deduction (CADE), no. 1632 in LNAI, pp. 287–291. Springer, Trento (1999)
84. Ong, C.L., Stewart, C.A.: A curry-howard foundation for functional computation with control. In: Conference Record of POPL'97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, Paris, France, 15-17 January 1997, pp. 215–227 (1997). DOI 10.1145/263699.263722. URL http://doi.acm.org/10.1145/263699.263722
85. Pareschi, R.: Type-driven natural language analysis. Ph.D. thesis, University of Edinburgh (1989)
86. Pareschi, R., Miller, D.: Extending definite clause grammars with scoping constructs. In: D.H.D. Warren, P. Szeredi (eds.) 1990 International Conference in Logic Programming, pp. 373–389. MIT Press (1990)
87. Paulson, L.C.: The foundation of a generic theorem prover. Journal of Automated Reasoning **5**, 363–397 (1989)
88. Pym, D.J., Harland, J.A.: The uniform proof-theoretic foundation of linear logic programming. J. of Logic and Computation **4**(2), 175–207 (1994)
89. Qi, X., Gacek, A., Holte, S., Nadathur, G., Snow, Z.: The Teyjus system – version 2 (2015). URL http://teyjus.cs.umn.edu/. http://teyjus.cs.umn.edu/
90. Robinson, J.A.: A machine-oriented logic based on the resolution principle. JACM **12**, 23–41 (1965)

91. Ruet, P., Fages, F.: Concurrent constraint programming and non-commutative logic. In: Proceedings of the 11^th Conference on Computer Science Logic, LNCS. Springer (1997)
92. Sangiorgi, D., Walker, D.: π-Calculus: A Theory of Mobile Processes. Cambridge University Press (2001)
93. Saraswat, V.: A brief introduction to linear concurrent constraint programming. Tech. rep., Xerox Palo Alto Research Center (1993). URL ftp://parcftp.xerox.com/pub/ccp/lcc/lcc-intro.dvi.Z
94. Sørensen, M.H., Urzyczyn, P.: Lectures on the Curry-Howard Isomorphism, *Studies in Logic*, vol. 149. Elsevier (2006)
95. Tammet, T.: A resolution theorem prover for intuitionistic logic. In: Automated Deduction — CADE-13, *LNCS*, vol. 1104, pp. 2–16 (1996)
96. Tarski, A.: Contributions to the theory of models. I. Indagationes Mathematicae **16**, 572–581 (1954)
97. Wadler, P.: Linear types can change the world! In: Programming Concepts and Methods, pp. 561–581. North Holland (1990)
98. Warren, D.H.D.: An abstract Prolog instruction set. Tech. Rep. 309, SRI International (1983)