A game semantics for proof search: Preliminary results

Dale Miller & Alexis Saurin
INRIA-Futurs & LIX

16 July 2005

Outline

1. A neutral approach to proof and refutation.
2. The noetherian Horn clause case.
3. Games for simple expressions.
4. Games for non-simple expressions.
5. Additive games and truth
6. Games for recursion.

Review: Horn clauses

The syntactic variable A denotes *atomic formulas*: that is, a formula with a predicate (a non-logical constant) as its head: the formulas $\bot$ and $\top$ and $t = s$ are *not* atomic formulas.

A *Horn goal* G is any formula generated by the grammar:

$$G ::= \top \mid \bot \mid t = s \mid A \mid G \wedge G \mid G \vee G \mid \exists x\, G.$$

A *Horn clause for* the predicate p is a formula

$$\forall x_1 \ldots \forall x_n [p(x_1, \ldots, x_n) \equiv G]$$

where $n \geq 0$, p is an n-ary predicate symbol, and the G, the *body*, is a Horn goal formula whose free variables are in $\{x_1, \ldots, x_n\}$.

A *Horn program* is a finite set $\mathcal{P}$ of Horn clauses all for distinct predicates.

Review: noetherian Horn clauses

Define $q \prec p$ to hold for two predicates if q appears in the body of the Horn clause for p.

$\mathcal{P}$ is *noetherian* if the transitive closure of $\prec$ is acyclic.

When $\mathcal{P}$ is noetherian, it can be rewritten to a logically equivalent logic program $\mathcal{P}'$ for which the relation $\prec$ is empty: that is, there are no atomic formulas in the body of clauses in $\mathcal{P}'$.

Repeatedly replace $\prec$-minimal predicates by their equivalent body.

Thus: in noetherian programs, atoms are not necessary.

Prolog and noetherian Horn clauses

Assume that the noetherian Horn clause program $\mathcal{P}$ is loaded into Prolog and we ask the query

```
?- G.
```

Prolog will respond by either reporting `yes` or `no`.

If `yes` then Prolog has a proof of G. Such a proof can be represented in "usual" sequent calculus (say, of, Gentzen).

If `no` then there is a proof of $\neg G$ in proof systems extended to deal with the *closed world assumption*: Clark's completion or more recent work on *definitions* and *fixed points* in proof theory (Schroeder-Heister & Hallnäs, Girard, and McDowell & Miller &Tiu).

Proof and refutation in one computation

This description of Prolog is a challenge to the conventional understanding of logic-as-proof-search paradigm (Miller, *et.al.*, in late 1980's).

Prolog did *one* computation which yielded a proof of G or a refutation of G (i.e., a proof of $\neg G$).

Proof search states that you must select first what you plan to prove and then proceed to prove that: i.e.,

$$\text{start with either} \quad \longrightarrow G \text{ or with} \quad \longrightarrow \neg G.$$

How can we *formalize* this neutral approach?

Can this behavior of Prolog be *extended* to richer logics?

A neutral approach to proof and refutation

Since a "neutral computation" could yield a proof of either $G_1 \wedge G_2$ or $\neg G_1 \vee \neg G_2$; or either $\exists x.G$ or $\forall x.\neg G$, we chose to compute with a new language of *neutral expressions*.

$$N ::= \mathbb{1} \mid N \times N \mid \mathbb{0} \mid N + N \mid \dot{p}\, t_1 \ldots t_n \mid \mathbf{Q} x N$$

Here, $\mathbb{1}$ and $\mathbb{0}$ are the units of $\times$ and $+$, respectively.

The expression $\dot{p}\, t_1 \ldots t_n$ is will correspond to the literal $pt_1 \ldots t_n$ or $\neg pt_1 \ldots t_n$.

The variable x in the expression $\mathbf{Q}x.N$ is bound in the usual sense.

First-order models, briefly

Let $\mathcal{M}$ be a *first-order model* in the usual sense.

- $|\mathcal{M}|$ denotes the domain of quantification of the model
- for every $c \in |\mathcal{M}|$ there is a *parameter* $\bar{c}$ in the language of the logic.
- An atomic formula $p(t_1, \ldots, t_n)$ is true if the n-tuple $\langle t_1^{\mathcal{M}}, \ldots, t_n^{\mathcal{M}} \rangle \in p^{\mathcal{M}}$.

Herbrand Models

Given a signature Σ, the model $\mathcal{H}_\Sigma$ is such that $|\mathcal{H}_\Sigma|$ is the set of closed terms built from Σ and in which the sole predicate that is interpreted is equality: $\mathcal{H}_\Sigma \models t = s$ if and only if t and s are identical closed terms.

Rewriting neutral expressions

Given a model $\mathcal{M}$ we describe a nondeterministic rewriting of multisets of neutral expressions.

$$\mathbb{1}, \Gamma \mapsto \Gamma \qquad N \times M, \Gamma \mapsto N, M, \Gamma$$

$$N + M, \Gamma \mapsto N, \Gamma \qquad N + M, \Gamma \mapsto M, \Gamma$$

$$p(t_1, ..., t_n), \Gamma \mapsto \Gamma, \qquad \text{if } \mathcal{M} \models p(t_1, ..., t_n)$$

$$\mathbf{Q}x.N, \Gamma \mapsto N[t/x], \Gamma, \quad \text{where } t \in |\mathcal{M}|$$

Let $\mapsto^*$ be the reflective and transitive closure of $\mapsto$.

Since expressions simplify, rewriting always terminates. Since the domain of quantification is infinite (all terms), rewriting can also be infinitely branching.

Main question: Given N, does $N \mapsto^* \{\}$?

Main proposition for Horn clauses over $\mathcal{H}_\Sigma$

Proposition. Let N be a neutral expression. If $N \mapsto^* \{\}$ then $\vdash [N]^+$. If N cannot be rewritten to $\{\}$ then $\vdash [N]^-$.

N	$[N]^+$	$[N]^-$
$\mathbb{0}$	0	$\top$
$\mathbb{1}$	1	$\bot$
$t \doteq s$	$t = s$	$\neg(t = s)$
$N_1 + N_2$	$[N_1]^+ \oplus [N_2]^+$	$[N_1]^- \,\&\, [N_2]^-$
$N_1 \times N_2$	$[N_1]^+ \otimes [N_2]^+$	$[N_1]^- \parr [N_2]^-$
$\mathbf{Q}x.N$	$\exists x.[N]^+$	$\forall x.[N]^-$

The range of $[\cdot]^+$ is a familiar linearization of *Horn goal* formulas.
The range of $[\cdot]^-$ is their negation.

Treatment of Equality

$$\frac{}{\vdash t = t} \qquad \frac{\vdash \Delta\theta}{\vdash \neg(t = s), \Delta}\,\dagger \qquad \frac{}{\vdash \neg(t = s), \Delta}\,\ddagger$$

The proviso † requires that t and s are unifiable and θ is their most general unifier ($\Delta\theta$ is the multiset resulting from applying θ to all formulas in Δ).

The proviso ‡ requires that t and s are not unifiable.

The free variables of a sequent are also called *eigenvariables*, which are introduced by the usual rule for $\forall$R.

Extending this neutral approach

Can we extend this neutral approach to proof and refutation beyond simple Horn goal formulas?

Proof search alternates between two phases.

- *asynchronous* phase where all inference rules are invertible. No choices need to be made.
- *synchronous* phase where inference rules require choices. A path through a proof must be made.

These two phases arise from dual aspects of the same logical connective.

So far, we only have one phase, with no alternation possible.

- *asynchronous* phase: all paths starting at N do not end in $\{\}$.
- *synchronous* phase: there is a path $N \mapsto^* \{\}$.

Adding the switch operator

Now add the *switch* operator to the language of neutral expressions.

$$N ::= \ldots \mid \updownarrow N.$$

Rewriting leaves switched expressions untouched.

Main question: Given N, does

$$N \mapsto^* \{\updownarrow N_1, \ldots, \updownarrow N_m\} = \updownarrow\{N_1, \ldots, N_m\}?$$

The motivation here:

(1) One player starts with her instructions N.

(2) She works on N in order to finish her "work", if possible.

(3) If she finishes successfully, she gives to the other player m instructions $N_1, \ldots, N_m$.

A class of *simple expressions* can be defined for which $m \leq 1$.

Games: Arenas, strategies, winning strategies

The pair $\langle P, \rho \rangle$ is an *arena*: P is a set of *positions* and ρ be a binary relationship on P that describes *moves*.

A *play* is a sequence $P_1.P_2.\cdots.P_n$ of ρ-related moves.

If σ is a set of plays then the set $\sigma/N = \{S \mid N.S \in \sigma\}$.

A *∀∃-strategy for* N is a prefixed closed set σ of plays such that $N \in \sigma$ and for all M such that $N \, \rho \, M$, the set σ/N is a ∃∀-strategy for M.

A *∃∀-strategy for* N is a prefixed closed set σ of plays such that $N \in \sigma$ and for at most one position M such that $N \, \rho \, M$, the set σ/N is a ∀∃-strategy for M.

A *winning ∀∃-strategy* is a ∀∃-strategy such that all its maximal sequences are of odd length. A *winning ∃∀-strategy* σ is a ∀∃-strategy such that all maximal sequences are of even length.

Games for simple expressions

Define $[\updownarrow N]^- = [N]^+$ and $[\updownarrow N]^+ = [N]^-$.

Let P be the set of neutral expressions. The move relation is defined as: $N\ \rho\ \mathbb{0}$ if $N \mapsto^* \{\}$ and $N\ \rho\ M$ if $N \mapsto^* \{\updownarrow M\}$.

Conjecture. Let N be a simple expression.
There is a winning $\forall\exists$-strategy for N if and only if $\vdash [N]^-$.
There is a winning $\exists\forall$-strategy for N if and only if $\vdash [N]^+$.

We have a number of examples supporting this Conjecture.

The Conjecture holds in the proposition case (when the model $\mathcal{M}$ is not relevant).

Example: finite sets

Encode $0, 1, 2, \ldots$ as terms $z, s(z), s(s(z)), \ldots$.

Let finite set $A = \{n_1, \ldots, n_k\}$ of natural numbers can be encoded as $A(x) = x \doteq n_1 + \cdots + x \doteq n_k$.

The expression $A(n)$ has a winning $\exists\forall$-strategy if and only if $n \in A$. In that case, $(n = n_1) \oplus \cdots \oplus (n = n_k)$ is provable.

The expression $A(n)$ has a winning $\forall\exists$-strategy if and only if $n \notin A$. In that case, $\neg(n = n_1) \,\&\, \cdots \,\&\, \neg(n = n_k)$ is provable.

If $A(x)$ and $B(x)$ encode two finite sets A and B, then the expressions $A(x) + B(x)$ and $A(x) \times B(x)$ encode in the intersection and union, respectively, of A and B.

Example: subset

The expression $\mathbf{Q}x.(A(x) \times \updownarrow B(x))$ encodes $A \subseteq B$.

Let P be the set $\{0, 2\}$ and let Q be the set $\{0, 1, 2\}$. The expression labeled $P \subseteq Q$, namely,

$$\mathbf{Q}x.([(x \doteq 0) + (x \doteq 2)] \times \updownarrow[(x \doteq 0) + (x \doteq 1) + (x \doteq 2)])$$

has a winning $\forall\exists$-strategy. Thus the following are provable.

$$\forall x.([\neg(x = 0) \mathbin{\&} \neg(x = 2)] \parr [(x = 0) \oplus (x = 1) \vee (x = 2)]).$$

$$\forall x.([(x = 0) \oplus (x = 2)] \multimap [(x = 0) \oplus (x = 1) \vee (x = 2)]).$$

The expression labeled $Q \subseteq P$, namely,

$$\mathbf{Q}x.([(x \doteq 0) + (x \doteq 1) + (x \doteq 2)] \times \updownarrow[(x \doteq 0) + (x \doteq 2)])$$

has a winning $\exists\forall$-strategy. Thus the following is provable:

$$\exists x.([(x = 0) \oplus (x = 1) \oplus (x = 2)] \otimes [\neg(x = 0) \mathbin{\&} \neg(x = 2)]).$$

Games for non-simple expressions

We do not know yet how to define games for general expressions.

Nor do we have any "computer science motivated" examples that indicate the need for non-simple expressions.

It is clear that such games cannot be determinate: that is, not all games will have either a winning $\forall\exists$-strategy or a winning $\exists\forall$-strategy.

For example, $\updownarrow\mathbb{1} \times \updownarrow\mathbb{1}$ should yield a game with *stuck states* since neither $1 \parr 1$ nor $\bot \otimes \bot$ are provable.

Additive Games and Truth

Hintikka showed that games can characterize truth in first-order logic.

Two players P and O play on the same formula:

- if that formula is a conjunction, then player P would choose one of the conjuncts;
- if is a universal quantifier, then player P would pick an instance;
- if the formulas is a disjunction, then player O picks a disjunct; and
- if the formula is an existential quantifier, play O picks an instance.

In our setting, such a game is purely additive: that is, the neutral expressions for such games contain no occurrences of $\times$ and $\mathbb{1}$.

Additive Games and Truth

Define two mappings, $f(\cdot)$ and $h(\cdot)$, from classical formulas in negation normal form (formulas where negations have only atomic scope) into additive neutral expressions.

$$
\begin{aligned}
f(B \wedge C) &= f(B) + f(C) & h(B \wedge C) &= \updownarrow f(B \wedge C) \\
f(B \vee C) &= \updownarrow h(B \vee C) & h(B \vee C) &= h(B) + h(C) \\
f(\top) &= \mathbb{0} & h(\top) &= \updownarrow f(\top) \\
f(\bot) &= \updownarrow h(\bot) & h(\bot) &= \mathbb{0} \\
f(\forall x.B) &= \mathbf{Q}x.f(B) & h(\forall x.B) &= \updownarrow f(\forall x.B) \\
f(\exists x.B) &= \updownarrow h(\exists x.B) & h(\exists x.B) &= \mathbf{Q}x.h(B) \\
f(\neg(p(t_1,\ldots,t_n))) &= \dot{p}(t_1,\ldots,t_n) & h(\neg A) &= \updownarrow f(A) \\
f(A) &= \updownarrow h(A) & h(p(t_1,\ldots,t_n)) &= \dot{p}(t_1,\ldots,t_n)
\end{aligned}
$$

Correctness of additive games with validity

Proposition. Let $\mathcal{M}$ be a model and let $f(E) = N$, where E is a closed first-order formula. The formula E is true in $\mathcal{M}$ if and only if there is a $\forall\exists$-win for N.

Proof. By simple induction over the structure of formulas.

Extending for recursion

Extend expressions with the fixed point constructors $\{\mathit{fix}_n\}_{n\geq 0}$. In

$$(\mathit{fix}_n \lambda P \lambda x_1 \ldots \lambda x_n . M)$$

the bound variable P is an n-ary recursive function. Extend $\mapsto$:

$$(\mathit{fix}_n F t_1 \ldots t_n), \Gamma \mapsto (F(\mathit{fix}_n F) t_1 \ldots t_n), \Gamma,$$

Extend the notions of winning strategies to infinite plays.

An *infinite play* is a *lose* for in a ∃∀-strategy while it is *win* for an ∀∃-strategy.

The *positive* translation of *fix* is the *least* fixed point operation μ; *negative* translation of *fix* is the *greatest* fixed point operation ν.

Example: less-than-or-equal

The logic program

```
leq(z,N).
leq(s(P),s(Q)) :- leq(P,Q).
```

can be written rather directly (using the Clark completion) as the expression

$$(\mathit{fix}_2\ \lambda leq \lambda n \lambda m[(n \doteq z) + \mathbf{Q}p\mathbf{Q}q.(n \doteq s(p) \times m \doteq s(q) \times leq(p,q))])$$

This expression, named L, has no $\updownarrow$ operator (it is just a Horn clause program).

$L(n,m)$ has a winning $\exists\forall$-strategy if and only if $n \leq m$.

$L(n,m)$ has a winning $\forall\exists$-strategy if and only if $n > m$.

Example: maximum

We can now define the maximum of a set of numbers. Let A be a non-empty set of numbers and let $A(n)$ be the expression encoding this set.

Let $maxA(n)$ be the following expression:

$$A(n) \times \updownarrow \mathbf{Q} m(A(m) \times \updownarrow L(m,n))$$

The expression $maxA(n)$ as a winning $\forall\exists$-strategy if and only if n is not in A or it is not the largest member of A. Similarly, $maxA(n)$ as a winning $\exists\forall$-strategy if and only if n is the largest member of A.

Example: bisimulation

Let $\delta \subseteq S \times \Lambda \times S$ be a finite transition on states S and labels Λ.

Encode this as the expression $\delta(x, y, z)$ given by

$$\sum_{(p,a,q)\in\delta} (x \doteq p \times y \doteq a \times z \doteq q).$$

Bisimulation between two states can be defined using the following recursive expression

$$\begin{aligned}(\mathit{fix}_2 \lambda bisim \lambda p \lambda q.\ & [\mathbf{Q} a \mathbf{Q} p'.\delta(p, a, p') \times \updownarrow \mathbf{Q} q'(\delta(q, a, q') \times \updownarrow bisim(p', q'))] \\ & + [\mathbf{Q} a \mathbf{Q} q'.\delta(q, a, q') \times \updownarrow \mathbf{Q} p'(\delta(p, a, p') \times \updownarrow bisim(p', q'))])\end{aligned}$$

If *Bisim* names the above expression and if p and q are two states (members of S), then the game for the expression $\mathit{Bisim}(p, q)$ is exactly the game usually used to describe bisimulation, eg., by C. Sterling.

Conclusions and Questions

- We have described a neutral approach to proof and refutation for an interesting and useful subset of logic (from the computer science point-of-view).
- Games and winning strategies provide a new way to look at proofs. This is not an approach to "full abstraction" for sequent proofs. We are hopeful for better "proof objects" than those.
- What is really going on with the multiplicatives?
- Can we extend this development to the modals (!, ?) of linear logic? To higher-order quantification?
- How does one implement the search for winning strategies using, say, unification?