15-819K: Logic Programming

Lecture 22

Hyperresolution

Frank Pfenning

November 14, 2006

In this lecture we lift the forward chaining calculus from the ground to the free variable case. The form of lifting required is quite different from the backward chaining calculus. For Horn logic, the result turns out to be hyperresolution.

22.1 Variables in Forward Chaining

Variables in backward chaining in the style of Prolog are placeholders for unknown terms. They are determined during unification, which occurs when the head of a program clause is compared to the current goal.

The same strategy does not seem appropriate for forward chaining. As a first example, consider

$$\forall x. Q(x) \vdash Q(a).$$

This violates the restrictions imposed for the last lecture, because x occurs in the head of $\forall x. Q(x)$ but not its body (which is empty).

We cannot focus on the right-hand side since no $Q(\_)$ is in the context. But we can focus on the assumption, $\forall x. Q(x)$. We have to guess a substitution term for x (which we will leave as a variable for now) and then blur focus.

$$\dfrac{\dfrac{\dfrac{\forall x. Q(x), Q(X) \vdash Q(a)}{\forall x. Q(x); Q(X) \ll Q(a)}\ \text{blur}L}{\forall x. Q(x); \forall x. Q(x) \ll Q(a)}\ \forall L}{\forall x. Q(x) \vdash Q(a)}\ \text{focus}L$$

Nothing within this focusing sequence gives us a clue to what X should be. If we now focus on the right, we can see in the second phase that X should be a in order to complete the proof.

But is the set $\forall x.\, Q(x), Q(X)$ actually saturated, or can we focus on $\forall x.\, Q(x)$ again? Since X is a placeholder, the proper interpretation of the sequent $\forall x.\, Q(x), Q(X) \vdash Q(a)$ should be:

There exists an X such that $\forall x.\, Q(x), Q(X) \vdash Q(a)$.

Now it could be we need two instances of the universal quantifier, so we might need to focus again on the left before anything else. An example of this is

$$\forall x.\, Q(x), Q(a) \supset Q(b) \supset Q'(c) \vdash Q'(c)$$

So a priori there is no local consideration to rule out focusing again on the left to obtain the context $\forall x.\, Q(x), Q(X), Q(Y)$, which is not redundant with $\forall x.\, Q(x), Q(X)$.

By extension of this argument we see that we cannot bound the number of times we need a universally quantified assumption. This means we can never definitively saturate the context. The existential interpretation of variables seems somehow incompatible with forward chaining and saturation.

22.2 Parametric Assumptions

Looking again at the deduction steps

$$\dfrac{\dfrac{\dfrac{\forall x.\, Q(x), Q(X) \vdash Q(a)}{\forall x.\, Q(x); Q(X) \ll Q(a)}\ \text{blur}L}{\forall x.\, Q(x); \forall x.\, Q(x) \ll Q(a)}\ \forall L}{\forall x.\, Q(x) \vdash Q(a)}\ \text{focus}L$$

we see that we lost a lot of information when blurring focus. We have actually obtain $Q(X)$ from the context without any restriction on what X is. In other words, we have really derived "*For any* X, $Q(X)$". When we added it back to the context, it became existentially quantified over the sequent: "*For some* $X, \ldots, Q(X) \vdash \ldots$".

It would make no sense to exploit this genericity of X by restoring the universal quantifier: we already know $\forall x.\, Q(x)$. Moreover, we would be introducing a connective during bottom-up reasoning which would violate all principles of proof search we have followed so far.

So we need to capture the fact that $Q(X)$ holds for any X in the form of a judgment. We write $\Delta \vdash Q$ where Δ is a context of variables. In a typed setting, Δ records the types of all the variables, but we ignore this slight generalization here. Now we have two forms of assumptions A and $\Delta \vdash Q$. We call the latter *parametric hypotheses* or *assumptions parametric in* Δ. There is no need to allow general parametric assumptions $\Delta \vdash A$, although research on contextual modal type theory suggests that this would be possible. The variables in Δ in a parametric assumption $\Delta \vdash Q$ should be considered bound variables with scope Q.

Parametric hypotheses are introduced in the blurring step.

$$\frac{\Gamma, (\Delta \vdash Q) \vdash Q' \quad \Delta = \mathrm{FV}(Q)}{\Gamma; Q \ll Q'}\ \mathrm{blur}L$$

We assume here that the sequent does not contain any free variables other than those in Q. Because for the moment we are only interested in forward chaining, this is a reasonable assumption. We discuss the issue of saturation below.

Now we consider the other rules of the focusing system, one by one, to see how to accomodate parametric hypotheses. We are restricting attention to the Horn fragment with only synchronous atoms.

We now rule out (non-parametric) assumptions Q and just allow $(\Delta \vdash Q)$. Closed assumptions Q are silently interpreted as $(\cdot \vdash Q)$.

The focus$L$ rule is as before: we can focus on any non-parametric A. By the syntactic restriction, this cannot be a Q, so we elide the side condition.

$$\frac{A \in \Gamma \quad \Gamma; A \ll Q}{\Gamma \vdash Q}\ \mathrm{focus}L$$

The implies$L$ rule also remains the same.

$$\frac{\Gamma \gg Q_1 \quad \Gamma; A_2 \ll Q}{\Gamma; Q_1 \supset A_2 \ll Q}\ {\supset}L$$

For the $\forall L$ rule we have to guess the substitution term t for x. This term t may contain some free variables that are abstracted in the blur step.

$$\frac{\Gamma; A(t/x) \ll Q}{\Gamma; \forall x.\, A \ll Q}\ \forall L$$

In the implementation t will be determined by unification, and we then abstract over the remaining free variables.

Focusing on the right is as before; the change appears in the identity rule

$$\frac{\Gamma \gg Q}{\Gamma \vdash Q}\ \text{focus}R \qquad \frac{(\Delta \vdash Q') \in \Gamma \quad Q'\theta = Q \quad \text{dom}(\theta) = \Delta}{\Gamma \gg Q}\ \text{id}L$$

Right focus on Q still fails if there is no appropriate $(\Delta \vdash Q')$ and θ.

22.3 Unification and Generalization

As a next step, we make the unification that happens during forward chaining fully explicit. This is the natural extension of matching during ground forward chaining discussed in the last lecture.

$$\frac{A \in \Gamma \quad \Gamma; A \ll Q}{\Gamma \vdash Q}\ \text{focus}L$$

$$\frac{\Gamma \gg Q_1 \mid \theta \quad \Gamma; A_2\theta \ll Q}{\Gamma; Q_1 \supset A_2 \ll Q} \supset L \qquad \frac{\Gamma; A(X/x) \ll Q \quad X \notin \text{FV}(\Gamma, A, Q)}{\Gamma; \forall x.\, A \ll Q}\ \forall L$$

$$\frac{(\Delta \vdash Q') \in \Gamma \quad \begin{array}{c}\rho \text{ renaming on } \Delta \\ Q'\rho \doteq Q \mid \theta\end{array}}{\Gamma \gg Q \mid \theta}\ \text{id}L \qquad \begin{array}{c}\text{no rule if no such } \Delta \vdash Q', \theta \\ \Gamma \gg Q \mid \theta\end{array}$$

$$\frac{\Gamma \gg Q \mid (\cdot)}{\Gamma \vdash Q}\ \text{focus}R \qquad \frac{\Gamma, (\Delta \vdash Q') \vdash Q \quad \Delta = \text{FV}(Q')}{\Gamma; Q' \ll Q}\ \text{blur}L$$

We do not permit free variables in Q for a global goal $\Gamma \vdash Q$. This may be reasonable at least on the Horn fragment if the renaming ρ always chooses fresh variables, since during forward chaining we never focus on the right except to complete the proof.

Reconsider an earlier example to see this system in action.

$$\forall x.\, Q(x), Q(a) \supset Q(b) \supset Q'(c) \vdash Q'(c)$$

We must focus on $\forall x.\, Q(x)$, which adds $y \vdash Q(y)$ to the context.

$$\forall x.\, Q(x), Q(a) \supset Q(b) \supset Q'(c), (y \vdash Q(y)) \vdash Q'(c)$$

Now we can focus on the second assumption, using substitutions a/y and b/y for the two premisses and adding $\cdot \vdash Q'(c)$ to the context. Now we can focus on the right to prove $Q'(c)$.

22.4 Saturation via Subsumption

In the ground forward chaining system of the last lecture we characterized saturation by enforcing that a blur step must add something new to the context Γ.

$$\frac{\Gamma, Q \vdash Q_0 \quad Q \notin \Gamma}{\Gamma; Q \ll Q_0}\ \text{blur}L'$$

We must update this to account for parametric hypotheses $\Delta \vdash Q$. One should think of this as standing for an infinite number of ground assumptions, $Q\theta$ where $\mathrm{dom}(\theta) = \Delta$ and $\mathrm{cod}(\theta) = \emptyset$.

We can say that $(\Delta \vdash Q)$ adds nothing new to the context if every instance $Q\theta$ is already an instance of a parametric assumption Q'. That is, for every θ there exists $(\Delta' \vdash Q') \in \Gamma$ and θ' such that $Q'\theta' = Q\theta$. A tractable criterion for this is *subsumption*. We way that $(\Delta' \vdash Q')$ *subsumes* $(\Delta \vdash Q)$ if there exists a substitution θ' with $\mathrm{dom}(\theta') = \Delta'$ and $\mathrm{cod}(\theta') \subseteq \Delta$ such that $Q'\theta' = Q$. Then every instance $Q\theta$ is also an instance of Q' since $Q\theta = (Q'\theta')\theta = Q'(\theta'\theta)$.

The new blur rule then is

$$\frac{\Gamma, (\Delta \vdash Q) \vdash Q_0 \quad \text{no } (\Delta' \vdash Q') \in \Gamma \text{ subsumes } (\Delta \vdash Q)}{\Gamma; Q \ll Q_0}\ \text{blur}L'$$

As an example, if we have a theory such as

$$\forall x.\, \mathsf{pos}(\mathsf{s}(x)), \forall y.\, \mathsf{pos}(y) \supset \mathsf{pos}(\mathsf{s}(y))$$

where, in fact, the second clause is redundant, we will saturate quickly. After one step we assume $(w \vdash \mathsf{pos}(\mathsf{s}(w)))$. Now focusing on the first assumption will fail by subsumption, and focusing on the second will also fail by subsumption. After unification during forward chaining we have to ask if $(u \vdash \mathsf{pos}(\mathsf{s}(\mathsf{s}(u))))$ is subsumed before adding it to the context. But it is by the previous assumption, instantiating $\mathsf{s}(u)/w$. Therefore the above theory saturates correctly after one step.

Under this definition of saturation it is possible to represent a number of decision procedures as saturating forward chaining search with subsumption. In general, however, we are straying more from logic programming into general theorem proving

22.5 Beyond Saturation

In many applications saturation may be possible, but suboptimal in the sense that we would like to short-circuit and succeed as soon as possible.

An example is the program for unification in the previous lecture. As soon as we have a contradiction to the assumption that two terms are unifiable, we would like to stop forward-chaining. We can achieve this by adding another synchronous connective to the logic: falsehood ($\bot$).

As a conclusion, if we are focused on $\bot$ we fail. So, just as in backward search, $\bot$ as a goal represents failure.

As an assumption it is asynchronous, so we can succeed when we encounter it.

$$\frac{}{\Gamma; \bot \ll C}\ \bot L$$

There is an implicit phase transition here, from focusing on $\bot$ to asynchronously decomposing $\bot$ (which immediately succeeds).

In the unification example, the uses of contra can be replaced by $\bot$, which sometimes permits early success when a contradiction has been derived.

This device is also used in theorem proving where we don't expect to saturate, but hope to derive $\bot$ from the negation of a conjecture. At this point we have left logic programming and are firmly in the real of general purpose theorem proving: we no longer try to implement algorithms, but try to search for a proof in a general way.

If we restrict ourselves to the Horn fragment (though allowing $\bot$), and every atom is synchronous then the strategy of forward chaining with free variables presented here is also known as *hyperresolution* in the theorem proving literature.

Once we have the possibility to succeed by creating a contradiction, it is no longer necessary to have a relevant right-hand side. For example, instead of proving $\Gamma \vdash Q$ we can prove $\Gamma, Q \supset \bot \vdash \bot$ entirely by forward chaining on the left, without ever considering the right-hand side. Most of the classical resolution literature and even early presentations of logic programming use this style of presentation. The proof is entirely by contradiction, and there is not even a "right-hand side" as such, just a database of facts and rules Γ.

22.6 Splitting

If we also allow disjunction in the heads of clauses, but continue to force all atoms to be synchronous, we can represent what is known in the theorem proving literature as *splitting*. Since disjunction is asynchronous on the left, we need a new judgment form $\Gamma; A \vdash C$ where A is broken down asynchronously. We transition into it when A is left asynchrounous, that is, Q, $\bot$, or $A_1 \vee A_2$. We give here the ground version.

$$\frac{\Gamma; A \vdash C \quad A = Q, \bot, A_1 \vee A_2}{\Gamma; A \ll C}\ \text{blur}L$$

$$\frac{}{\Gamma; \bot \vdash C}\ \bot L \qquad \frac{\Gamma; A_1 \vdash C \quad \Gamma; A_2 \vdash C}{\Gamma; A_1 \vee A_2 \vdash C}\ \vee L$$

$$\frac{\Gamma, A \vdash C \quad A \neq \bot, A_1 \vee A_2}{\Gamma; A \vdash C}$$

Unfortunately, this extension interacts poorly with free variables and parametric hypotheses. If there is a variable shared between A_1 and A_2 in the $\vee L$ rule, then it must be consistently instantiated on both sides and may not be abstracted. In the theorem proving context the rule is therefore restricted to cases where A_1 and A_2 share no free variables, which leaves the calculus complete for classical logic. Here, in intuitionistic logic, such a restriction would be incomplete.

When we move a formula A into Γ during decomposition of left asynchronous operators, we need to be able to abstract over its free variables even when the formula A is not atomic, further complicating the system.

To handle such situations we might allow existentially interpreted free variables in the context, and abstract only over those that are not free elsewhere in the sequent. However, then both subsumption and saturation become questionable again. It seems more research is required to design a larger fragment of intuitionistic logic that is amenable to a forward chaining operational semantics with reasonable saturation and subsumption behavior.

22.7 Historical Notes

Endowing assumptions with local contexts is a characteristic of contextual modal type theory [3] and the proof theory of the Nabla quantifier (∇) [2]. The former matches the use here, but is somewhat more general. The latter interprets the locally quantified variables as names subject to α conversion but cannot be instantiated by arbitrary terms.

There are a number of papers about using saturating hyperresolution as a decision procedure. A tutorial exposition and further references can be found in a chapter in the *Handbook of Automated Reasoning* [1].

22.8 References

[1] Christian Fermüller, Alexander Leitsch, Ullrich Hustadt, and Tanel Tammet. Resolution decision procedures. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume 2, chapter 25, pages 1791–1849. Elsevier Science and MIT Press, 2001.

[2] Dale Miller and Alwen Tiu. A proof theory for generic judgments. *ACM Transactions on Computational Logic*, 6(4):749–783, October 2005.

[3] Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. Submitted, September 2005.