

15-424/15-624 Recitation 8
Proof intuitions and KeYmaera specifics

1. A diversion: what if there's no sequent?

Take the proof rule for $[\cup]$.

$$[\cup] \frac{[\alpha]\phi \wedge [\beta]\phi}{[\alpha \cup \beta]\phi}$$

Where did that darned sequent (\vdash) go? Do we need it? What does it mean if it's not there?! The answer to these questions *and more* in the next paragraph, so keep reading!

If there's no sequent, then the understanding is that the rule can be applied to either the antecedent or the succedent of a sequent! Notice how the proof will end up looking different depending on which side it is.

In the antecedent, we'll get to use \wedge_L and simply add both sides of the conjunction to the assumptions. The proof doesn't branch. In the succedent, we now have to prove *two* things, meaning the proof will branch.

2. A simpler unwind rule

How can we deal with $\langle \alpha^* \rangle \phi$? Of course, with a proof rule, but which one! There's a very general one in the lecture notes, but your TAs are enormous procrastinators, and prefer simpler rules.

The meaning of the formula is clear enough: is there an execution of α^* such that ϕ holds at the end. In other words, can we find a number of iterations n of the program α which leads us to stop at a state where ϕ holds?

$$\langle *^n \rangle \frac{\langle \alpha^n \rangle \phi}{\langle \alpha^* \rangle \phi}$$

This is actually a family of proof rules! There's one for each $n \in \mathbb{N}$. If we can get away with simpler rules by using these "families" instead of a single rule, why shouldn't we?

Well, we dare you to try to prove a non-trivial property using this rule! Can you figure out the exact n for which this will work? When dealing with complex systems, the clear answer to that is going to be a flat "no". In other words, we're simplifying the rule itself, but unfortunately we are giving the prover (whether that's you or KeYmaera) the burden of figuring out the hard choice. And that's not very helpful, which is why the other rule is used.

Now, because we procrastinated and came up with a new proof rule, that comes back and bites us in the ***. We have to prove its soundness!

In this instance, the property is mainstream and cool because it contains an ODE. So the substitution rule can't be applied, and the above proof is *unsound*. The problem that the assignment is only meaningful at time $t = 0$ of the ODE, but the substitution is making changes *throughout* the execution of the ODE.

A similar phenomenon can be seen in the following:

$$[:=] \frac{[(x := 4)^*]\phi}{[x := 3][(x := x + 1)^*]\phi}$$

This clearly changes the behaviour of the program in a fundamental way. The cool insight is that instead of applying the substitution immediately, we can *delay it*, much like KeYmaera does. After we get rid of the programs, it's then possible to perform the substitution!

$$\begin{array}{l} \text{update simpl.} \\ \text{ODE solve} \end{array} \frac{[x := x - cvt + \frac{1}{2}At^2, v := -cv + at]\phi}{\frac{\{v := -cv\}[x := x + vt + \frac{1}{2}At^2, v := v + at]\phi}{[:=] \frac{\{v := -cv\}[x' = v, v' = A]\phi}{[v := -cv][x' = v, v' = A]\phi}}}$$

Notice that after applying ODE solve, all the occurrences of v and x refer to their *original* values, which is why it's become sound to perform the substitution. In loops and ODEs, an occurrence of a variable refers to their value at whatever iteration or time of the loop or ODE it refers to.

5. ind loop vs ind loop (global)

There are two variants of loop rules that can be used in KeYmaera. One of them allows you to keep your original assumptions, the other one does not. They do so by using the \forall^α construct, which universally quantifies all of the *bound* variables of program α .

What this essentially does is it “hides” the original value of the variable. Thus, it allows you to keep all the information that you had from before, since the relevant parts have been “hidden” through quantification.

You are encouraged to look at the attached `inv.key` program, and prove it with the invariant $x > 0$, using both `ind` and `ind (global)` rules. You'll notice that, to keep things sound, one of them hides information about the non-relevant y , whereas as the other does not.

When previous information about non-bound variables is hidden, you must include it in the invariant (since it doesn't change anyway). When it isn't, you have perhaps too much arithmetic lying around, and that might put a strain on QE. But you can always hide that information later.