15-424/15-624/15-824 Recitation 6
Differential Invariants and Differential Cuts

Reminder: The first exam is next Thursday, so today we're going to focus on material that's important for the exam. Specifically, we're going to focus on differential invariants because they're new, but also important for the exam. In contrast, we will not talk about differential ghosts yet until after the exam. They're cool and useful, but it's more important to focus on understanding differential invariants well before moving on to ghosts.

Today's example model is a frictionless box sliding down a ramp. The physics are non-obvious, so let's start with the physics before doing the proofs.

1 Modeling a Box on a Ramp

In Newtonian mechanics, if we have a frictionless box on a ramp, there are two forces (or equivalently, accelerations) that we have to model. The first force is gravity, which pulls the box straight down. If that's the only force we had, then the box would just fall through the ramp, so the ramp has to push back at the box, but it can only push back along the *normal vector*. How strong does this *normal force*(or rather acceleration due to normal force) have to be before it stops the box from falling through the ramp? It needs to be just strong enough that the overall acceleration will be parallel to the ramp (so we slide *along* the ramp instead), which is a fancy way of saying $g \cdot \cos(\theta)$, where g is acceleration due to gravity and θ is the angle between the normal vector and the *x-axis*. We can break this down into an x-component and y-component of acceleration:

$$v_x' = g \cdot \cos(\theta)\sin(\theta), v_y' = g \cdot (\cos^2(\theta) - 1)$$

Alarm bells should be going off in your head right now. Trigonometric functions like $\cos(\theta)$ are a major pain to prove things about, so much that KeYmaera X doesn't support them. What are we going to do? The trick is that we have an easy definition for $\cos(\theta)$ and $\sin(\theta)$ *if* we already know $\tan(\theta)$. And if you think about it, $\tan(\theta)$ is really simple here: that's just the slope of the ramp. Let's make the slope a parameter to our model and call it m. Then recall from trigonometry that we can compute:

$$\cos(\theta) = \frac{1}{\sqrt{1+m^2}}, \sin(\theta) = \frac{m}{\sqrt{1+m^2}}$$

We should still be a little careful: We formally haven't defined square roots as part of $\mathsf{d}\mathcal{L}$, but KeYmaera X is nice to us and supports them anyway. Those formulas are a bit

complicated, but they're never going to change because the slope is constant, so let's stick them in variables so we can reuse them:

$$cosTh := \frac{1}{\sqrt{1+m^2}}; sinTh := \frac{m}{\sqrt{1+m^2}}$$

We can stick this together to get a hybrid program:

$$\alpha \equiv cosTh := \frac{1}{\sqrt{1+m^2}}; sinTh := \frac{m}{\sqrt{1+m^2}}; \{x' = v_x, v_x' = g{\cdot}cosTh{\cdot}sinTh, , y' = v_y, v_y' = g{\cdot}(cosTh^2{-}1)\}$$

Now that we have a model, we should do what we do every time we come up with a cool new model: Prove stuff about it, this time using differential invariants!

2 Proving Invariant Equations

Let's warm up with a familiar invariant: proving that the box always stays on the ramp. To make things simple, let's say the ramp passes through the origin, in which case the ramp is just the line $y = m \cdot x$. Last week we had to dream up a new invariant in order to prove we stayed on the line, but now we have a new, better version of DI that can prove arbitrary equalities:

$$\text{DI}_{=}\ \frac{\Gamma \vdash \theta_1 = \theta_2, \Delta \qquad \Gamma_{\texttt{const}}, Q \vdash [x' := \theta](\theta_1)' = (\theta_2)'}{\Gamma \vdash [x' = \theta \& Q](\theta_1 = \theta_2), \Delta}$$

What this says is that if two terms θ_1, θ_2 are equal at the start of an ODE, they remain equal so long as their differentials $(\theta_1)'$ and $(\theta_2)'$ are equal. It's important to pay attention to the contexts here: We can't assume the invariant $\theta_1 = \theta_2$ holds when proving $(\theta_1)' = (\theta_2)$. If we want to really keep things simple, we should just say that you can't keep around *any* assumptions except what's in the domain constraint. In practice, that a little unwieldy, so instead we say you can only keep around formulas which "obviously" never change, where "obvious" has a precise technical definition: formulas that don't mention any variables which are modified by the ODE. Our notation for this is $\Gamma_{\texttt{const}}$. The intuition for this proof rule can be described graphically:

Let's go ahead and try using this proof rule to show that we stay on the ramp, using the following definitions to make the proof easier to read:[1]

$$\Gamma \equiv x = 0, y = 0, v_x = 0, v_y = 0, cosTh = \frac{1}{\sqrt{1+m^2}}, sinTh = \frac{m}{\sqrt{1+m^2}}$$

[1]Man, what a handy technique to reduce the time spent writing when you're writing a proof in a time-pressured environment!

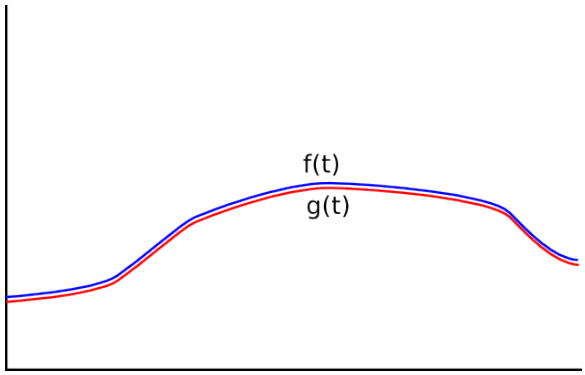

Figure 1: Invariant for $\theta_1 = \theta_2$: $(\theta_1)' = (\theta_2)'$

$$\Gamma_{\texttt{const}} \equiv cosTh = \frac{1}{\sqrt{1+m^2}}, sinTh = \frac{m}{\sqrt{1+m^2}}$$
$$\beta \equiv \{x' = v_x, v_x' = g \cdot cosTh \cdot sinTh, y' = v_y, v_y' = g \cdot (cosTh^2 - 1)\}$$

$$\text{DI}_= \frac{\Gamma \vdash y = m \cdot x \qquad \text{derive}\,\dfrac{[:=]\,\dfrac{???\,\dfrac{???}{\Gamma_{\texttt{const}} \vdash v_y = -m \cdot v_x}}{\Gamma_{\texttt{const}} \vdash [\beta]y' = -m \cdot x'}}{\Gamma_{\texttt{const}} \vdash [\beta](y)' = (-m \cdot x)'}}{\Gamma \vdash [\beta]y = -m \cdot x}$$

We can follow the strategy for doing DI proofs, but eventually we get stuck. If we look at the one remaining branch, it's asking us "does the velocity vector always stay along the ramp?" That sounds like something that should actually be true, though, so let's try proving it.

$$\text{DI}_= \frac{\mathbb{R}\,\dfrac{*}{\Gamma \vdash v_y = m \cdot v_x} \qquad \text{derive}\,\dfrac{[:=]\,\dfrac{\text{algebra}\,\dfrac{\mathbb{R}\,\dfrac{*}{\Gamma_{\texttt{const}} \vdash g \cdot (\frac{-m^2}{m^2+1}) = -m \cdot g\frac{m}{m^2=1}}}{\Gamma_{\texttt{const}} \vdash g \cdot (\frac{1}{m^2+1} - 1) = -m \cdot g\frac{m}{m^2=1}}}{\Gamma_{\texttt{const}} \vdash [\beta]v_y' = -m \cdot v_x'}}{\Gamma_{\texttt{const}} \vdash [\beta](v_y)' = (-m \cdot v_x)'}}{\Gamma \vdash [\beta]v_y = -m \cdot v_x}$$

Are we done? Can we just plug this into the hole in our other proof? Nope! Why? Because this proof used assumptions from Γ, such as the initial values $v_x = 0, v_y = 0$. Remember we lose the (non-constant parts of the) context Γ in the middle of a differential invariant proof, so we can't copy-paste this proof into ???. We are *almost* done, though. We proved a lemma about our ODE (velocity vector always follows the slope) and now we want to use that to prove a theorem about the same ODE (the position always follows the slope). This is exactly what a *Differential Cut* allows us to do:

$$\frac{\Gamma \vdash [x' = \theta \& Q]R, \Delta \qquad \Gamma \vdash [x' = \theta \& (Q \wedge R)]P, \Delta}{\Gamma \vdash [x' = \theta \& Q]P, \Delta}$$

Differential cuts are named after the *cut* rule from sequent calculus, and follow a similar intuition:

$$\frac{\Gamma \vdash P, \Delta \qquad \Gamma, P \vdash, \Delta}{\Gamma \vdash \Delta}$$

If we really wanted to prove some formula $Q \in \Delta$, a perfectly valid way to do that is to first go and prove a lemma P and use P to help finish proving Q. Similarly, if we want to prove that some formula P is a differential invariant, we can first prove some other invariant R and add it to the domain constraint, where it can help us prove P. In this sense, domain constraints work kind of like contexts Γ for a differential equation.

If we use a differential cut, now when we get to point ??? we'll just have to prove:

$$v_y = -m \cdot v_x, \Gamma_{\texttt{const}} \vdash [\beta] v_y = -m \cdot v_x$$

Wow, the domain constraint is exactly what we wanted to prove, so we can prove this using by `close`!

$$[:=] \frac{\text{id} \dfrac{*}{v_y = -m \cdot v_x, \Gamma_{\texttt{const}} \vdash v_y = -m \cdot v_x}}{v_y = -m \cdot v_x, \Gamma_{\texttt{const}} \vdash [\beta] v_y = -m \cdot v_x}$$

More generally, this illustrates our methodology for proving differential invariants:

1. Try proving some formula is an invariant using DI
2. If the proof fails, look at where it failed, use the failure to come up with a new invariant ϕ. Cut in ϕ with differential cut and repeat.
3. When you have enough invariants, use differential weakening to finish the proof.

And in fact, if you're starting a proof on paper and you suspect you're going to need a differential cut, it might save you time to just do a cut right away and name the cut formula R. Further along the proof you'll eventually discover what R should be, then you can come back and define + prove R.

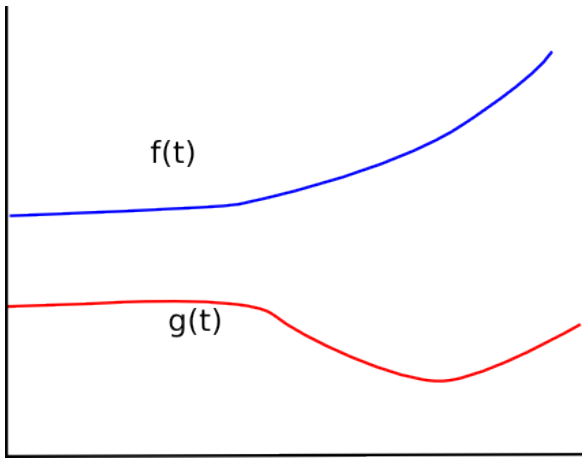

Figure 2: Invariant for $\theta_1 \geq \theta_2$: $(\theta_1)' \geq (\theta_2)'$

3 Proving Invariant Inequalities

We also saw in lecture that we can prove inequalities by invariant as well. The intuition is not so different from equations: If $x \geq y$ and x, y only keep getting further apart, then it will always be the case that $x \geq y$.

$$\frac{\Gamma \vdash \theta_1 \geq \theta_2, \Delta \qquad \Gamma_{\texttt{const}}, Q \vdash [x' := \theta](\theta_1)' \geq (\theta_2)'}{\Gamma \vdash [x' = \theta \& Q](\theta_1 \geq \theta_2), \Delta}$$

Let's use this to prove more fun things about boxes. It turns out our friend Nathan has a box too, and Nathan is really competitive, so he wants to have a box race! Luckily, our ramp has a steeper slope than his, which means we're going to win the race, but just to be safe let's prove that we win the race. First we update the model so we have two boxes on two ramps:

$$\begin{aligned}
\alpha' \equiv \\
\{v'_{x,1} =& g \cdot cosTh_1 \cdot sinTh_1, & v'_{y,1} &= g \cdot (cosTh_1^2 - 1), \\
v'_{x,2} =& g \cdot cosTh_2 \cdot sinTh_2, & v'_{y,2} &= g \cdot (cosTh_2^2 - 1), \\
x'_1 =& v_{x,1} \qquad x'_2 = v_{x,2} \qquad y'_1 = v_{y,1} \qquad y'_2 = v_{y,2}\}
\end{aligned}$$

and we add some assumptions: we both start out at a stop and our ramp is steeper:

$$\Gamma' \equiv x_1 = 0, y_1 = 0, x_2 = 0, y_2 = 0, v_{x,1} = 0, v_{y,1} = 0, v_{x,2} = 0, v_{y,2} = 0, m_1 > m_2$$

$$cosTh_1 = \frac{1}{\sqrt{1+m_1^2}}, sinTh_1 = \frac{m_1}{\sqrt{1+m_1^2}}, cosTh_2 = \frac{1}{\sqrt{1+m_2^2}}, sinTh_2 = \frac{m_2}{\sqrt{1+m_2^2}}$$

The proof follows a similar structure: First we'll need to prove a lemma about the velocities and then we can prove the theorem about the positions. What's the lemma in this case? We need to know that we're always moving faster than the other box, which is intuitively relevant to whether we're winning the race.

$$\mathrm{DC}\;\dfrac{\mathrm{DI}_{\geq}\;\dfrac{\mathbb{R}\;\dfrac{*}{\Gamma \vdash y_2 \geq y_1} \qquad [:=]\;\dfrac{\mathrm{id}\;\dfrac{*}{\Gamma_{\mathtt{const}}, v_{y,2} \geq v_{y,1} \vdash v_{y,2} \geq v_{y,1}}}{\Gamma_{\mathtt{const}}, v_{y,2} \geq v_{y,1} \vdash [\beta'](y_2)' \geq (y_1)'}}{\Gamma \vdash [\alpha' \wedge v_{y,2} \geq v_{y,1}] y_2 \geq y_1} \qquad \mathrm{DI}_{\geq}\;\dfrac{\mathbb{R}\;\dfrac{*}{\Gamma \vdash (v_{y,2} \geq v_{y,1})} \qquad \text{eq. rewriting}\;\dfrac{[:=]\;\dfrac{\mathbb{R}\;\dfrac{*}{\Gamma_{\mathtt{const}} \vdash g \cdot \frac{m_2}{1+m_2^2} \geq g \cdot \frac{m_1}{1+m_1^2}}}{\Gamma_{\mathtt{const}} \vdash g \cdot cosTh_2 \cdot sinTh_2 \geq g \cdot cosTh_1 \cdot sinTh_1}}{\Gamma_{\mathtt{const}} \vdash [\beta']((v_{y,2})' \geq (v_{y,1})')}}{\Gamma \vdash [\alpha'](v_{y,2} \geq v_{y,1})}}{\Gamma \vdash [\alpha'](y_2 \geq y_1)}$$

2

4 Choosing Cuts + Invariants

There's an interesting pattern in both of the previous proofs: First we proved something about v and used it to prove something about x or y. This is more than just a cute observation: In any proof we ever do, the order in which we cut+prove differential invariants is going to follow the dependencies between the variables. Because x depends on v, we have to prove something about v first. If multiple variables depend on each other mutually, then we would prove an invariant about them simultaneously (think circular motion dynamics).

5 Differential Formulas: Arbitrary Invariants

Both of the DI rules we've used so far are just special cases of the mother-of-all-DI-rules:

$$\frac{\Gamma \vdash P, \Delta \qquad \Gamma_{\texttt{const}}, Q \vdash [x' := \theta](P)'}{\Gamma \vdash [x' = \theta \& Q]P, \Delta}$$

We can prove any formula P using differential induction by proving that its differential $(P)'$ is always true and P is true initially. But what does $(P)'$ actually *mean*? Unlike all the other formulas in $\mathsf{d}\mathcal{L}$, we haven't bothered to tell you the semantics of $(P)'$ yet. And the idea of differentials for formulas is *weird*. I mean, formulas are either true or false, so how the heck would they ... change ... continuously?

It's actually quite difficult to come up with a beautiful semantics for $(P)'$ (read: nobody has done that yet). But observe that the rule DI above is the *only* place we ever use the $(P)'$ operator in $\mathsf{d}\mathcal{L}$, so the following informal semantics is sufficient (well, except we wish it was formal):

$[[(P)']] = [[\phi]]$For any formula ϕ so long as $\vdash [x' = \theta \& Q]\phi$ and $\Gamma \vdash P, \Delta$ give us $\Gamma \vdash [x' = \theta \& Q]P, \Delta$

3

[2]The astute reader may notice this is a different proof from what was done in recitation. This is because the property proved in recitation (as pointed out by an astute student) is false. But what we can prove is that we always win the y-coordinate race (meaning we always have a more-negative y-coordinate than the opponent. Luckily the structure of the proof is exactly the same; the only part that was wrong is the arithmetic that I brushed over quickly in recitation anyway.)

[3]Interestingly enough, there is actually a nice way to express this definition in logic, and your TA almost did that. If you're curious, ask about Hilbert's Epsilon Operator.

In other words, it doesn't really matter *which* invariant $(P)'$ gives us, as long as it happens to be a sound invariant. So let's look at the different cases for P and figure out a sound invariant in each case:

$$
\begin{aligned}
(\theta_1 = \theta_2)' &= (\theta_1)' = (\theta_2)' \\
(\theta_1 \geq \theta_2)' &= (\theta_1)' \geq (\theta_2)' \\
(\theta_1 > \theta_2)' &= (\theta_1)' \geq (\theta_2)' \\
(P \wedge Q)' &= (P)' \wedge (Q)' \\
(P \vee Q)' &= (P)' \wedge (Q)' \\
(\theta_1 \neq \theta_2)' &= (\theta_1)' = (\theta_2)'
\end{aligned}
$$

We've actually already seen why the cases for $=$ and $\geq$ are sound — those are the previous two version of the DI rule. The case for $\theta_1 > \theta_2$ is very similar, but notice we can get away with a nonstrict inequality $(\theta_1) \geq (\theta_2)'$ because if the derivatives are exactly the same, then θ_1 is still greater than θ_2 by just as much as before.

What about $P \wedge Q$? Well in general if we want to prove $[\alpha](P \wedge Q)$ (it doesn't even matter whether α is an ODE), one way we can do that is by proving both $[\alpha]P$ and $[\alpha]Q$ and sticking them together with the $[]\wedge$ axiom. And how would we prove $[\alpha]P$ and $[\alpha]Q$ in this case? By calling DI twice, which means proving $[\alpha](P)'$, $\vdash [\alpha](Q)'$, $\Gamma \vdash P$, and $\Gamma \vdash Q$. So $(P)' \wedge (Q)'$ worked as an invariant.

What about $P \vee Q$? It only seems natural that $(P \vee Q) = (P)' \vee (Q)'$. Unfortunately that would also be *wrong*. As before, let's imagine how we might try to "implement" the $\vee$ case of DI. If we know that $P \vee Q$ holds true initially, we could case on which one was true: If P is true then prove $[\alpha]P$ by DI, else if Q is true prove $[\alpha]Q$ by DI. There's only one way you can do that, though: when you know *both* $[\alpha](P)'$ and $[\alpha](Q)'$! If you only required $[\alpha](P)' \vee [\alpha](Q)'$ then you might end up in the case where all you have is $P \wedge [\alpha](Q)'$, in which you get stuck.

What about $\theta_1 \neq \theta_2$? You can actually derive this using the rules above. You can also prove it directly: If we had $(\theta_1)' \neq (\theta_2)'$ we have no idea what they'll look like after the ODE finishes, but if $(\theta_1)' = (\theta_2)'$ then the difference $\theta_1 - \theta_2$ will be constant throughout the ODE, so $\theta_1 \neq \theta_2$ is preserved.

You might be disappointed by the rules for $\vee$ and $\neq$ because it feels like they throw away some information (i.e. they are incomplete). If you felt this way, you were right! Sometimes you might find that using differential induction on some formula ϕ won't prove it, but if you rewrite it to some equivalent ψ, DI will work. The reason is that DI uses the *syntax* of ϕ to pick an invariant, not the *semantics*. Even when syntactically different ϕ and ψ are *semantically equivalent*, their invariants may be semantically *different*! Watch out for this if you get stuck in a DI proof.

The good news is that you're totally allowed to replace ϕ with ψ when you can prove that $\phi \leftrightarrow \psi$ is valid. When DI is combined with our rewriting superpowers, they *are* complete.

Case Study: 3D Lotka-Volterra

In fact, here's a concrete example of a system where you need to rewrite the formula before doing DI. The following predator/prey model describes the behavior of the biomasses x, y and z of three distinct species. That is, it determines how the population of each species evolves over time (when they get eaten by each other and whatnot). We want to prove that none of the three involved species will disappear: that is we reach an equilibrium cycle.

```
ProgramVariables.
  R x.
  R y.
  R z.
End.
Problem.
  x != 0 & y != 0 & z != 0
  ->
  [
    {x' = x*(y-z),
     y' = y*(z-x),
     z' = z*(x-y)
    }
  ](x != 0 & y != 0 & z != 0)
```

1. Apply a DI first (with the postcondition as differential invariant). Observe that the proof does not close because the condition asks about separate properties for x, y and z.
2. Apply a DC with $xyz \neq 0$ (which is equivalent to the post-condition, but links explicitly the involved variables).
3. Close the proof by a DI and DW.

Quiz

1. Can you prove that $y > 0 \land x < 0 \rightarrow [x' = x, y' = y]x \neq y$? Explain why or why not.
2. Can you prove $x < x_o \rightarrow [a := \frac{v^2}{2(x-x_o)}; \{x' = v, v' = a, v \geq 0\}]x \leq x_o$ using DI instead of ODE (solving the differential equation) ? Write down your DI.
3. Try proving the `.kyx` files in the recitation `.zip` file.