Lecture notes for PM 464/764 – Week Nine

David McKinnon
Department of Pure Mathematics
University of Waterloo

Spring 2021

1 DVRs and their reign over curves

It's handy to know a few things about DVRs. Like this, for example:

Theorem 1.1. *Let D be a DVR, $M = (t)$ its maximal ideal, K its fraction field. Then every ideal of D is principal, generated by t^n for some integer n, and every nonzero element x of K can be written as $x = ut^n$ for some integer n and some unit $u \in D^*$ of D.*

The integer n is called the **order of vanishing** *of x, written $ord(x)$.*

Proof: We will start by showing that any nonzero element $y \in D$ can be written as $y = ut^n$ for some unit u and some integer n.

If y is a unit, then we're clearly already done. (Set $n = 0$ and $u = y$.)

If y is not a unit, then $y \in M$, so $y = y_1 t$ for some $y_1 \in D$.

If y_1 is a unit, then we're done. If not, then we can write $y_1 = y_2 t$ for some $y_2 \in D$.

Keep going in this way until y_i is a unit.

"Why does this happen?" I hear you ask. (I have very good hearing. I don't just hear what people actually say. I also hear what they *should* say.)

Certainly we have an inclusion of ideals $(y_i) \subset (y_{i+1})$ for all i. That means we have a chain

$$(y_1) \subset (y_2) \subset \dots$$

which, because D is Noetherian, must stabilize in a finite number of steps. That is, there is some i when $(y_i) = (y_j)$ for all $j > i$.

But $y_i = y_{i+1}t$, so if $(y_i) = (y_{i+1})$, then t must be a unit ... which it isn't. Therefore, it can't actually happen that $(y_i) = (y_{i+1})$, meaning that at some stage, y_i must have been a unit.

Great! Because if y_i is a unit, then we can set $u = y_i$, giving $y = ut^i$, as desired.

Now let I be any ideal of D. Since D is Noetherian, I is generated by a finite number of elements, $\{z_1, \dots, z_r\}$. We just proved that for each i, we have

$$z_i = u_i t^{n_i}$$

for some unit u_i and some (non-negative) integer n_i. Since $ut^n \in (ut^m)$ if and only if $n \geq m$, it follows that I is generated by any z_i for which n_i is minimum. In particular, I is principal, with $I = (u_i t^{n_i}) = (t^{n_i})$, as advertised.

Now let $x \in K$ be any nonzero element. We can write $x = y/z$ for $y, z \in D$, and we can further write $y = ut^n$ and $z = vt^m$ for units u and v. Which immediately gives

$$x = \frac{ut^n}{vt^m} = wt^{n-m}$$

where $w = uv^{-1}$ is a unit. Woo! ♣

All this turns out to have some pretty amazing consequences. For example:

Theorem 1.2. *Let $C \subset \mathbb{P}^m$ be a smooth curve, and $\phi\colon C \to \mathbb{P}^n$ a rational map from C to projective space. Then ϕ is a morphism. (That is, ϕ is defined at every point of C.)*

True fact.

Proof: Let P be a point of C. We want to show that ϕ is defined at P.

Write $\phi = [\phi_0 : \ldots : \phi_n]$, where the ϕ_i are homogeneous polynomials of the same degree. If $\phi_i(P) \neq 0$ for some i, then ϕ is defined at P and we're all done.

Now, it's certainly true that at least one of the ϕ_i isn't the zero polynomial (although some of them could be), otherwise ϕ isn't defined anywhere. So, without loss of generality, let's say that ϕ_0 isn't the zero polynomial. That means that $\phi_1/\phi_0, \ldots, \phi_n/\phi_0$ are all rational functions on C.

But $\mathcal{O}_P(C)$ is a DVR. Which means that for each i, we can write

$$\frac{\phi_i}{\phi_0} = u_i t^{m_i}$$

where u_i is a unit of $\mathcal{O}_P(C)$, t is a uniformizer for $\mathcal{O}_P(C)$, and m_i is an integer (possibly negative).

This lets us write

$$\begin{aligned}\phi &= [\phi_0 : \ldots : \phi_n] \\ &= [1 : \phi_1/\phi_0 : \ldots : \phi_n/\phi_0] \\ &= [1 : u_1 t^{m_1} : \ldots : u_n t^{m_n}]\end{aligned}$$

Remember that all these u_i and t and stuff are just rational functions, modulo the ideal of C. Nothing mysterious, really.

Anyway, this new representation still might not be defined at P, because some of the powers of t might be negative. (But the rest is fine – a unit is defined at P for sure, and so is a uniformizer.)

The power of projective space will still save us, though. Say that $u_i t^{m_i}$ has the *most* negative m_i. (And assume that it's actually negative, because otherwise we're already done.) Simply multiply all the coordinates by t^{-m_i}, and you get a new representation with no negative powers of t. And you don't have to worry that all the coordinates will be zero at P, because the ith coordinate turns out to be just u_i, which is nonzero there (it's a unit!).

So ϕ is defined at every point of C, just like I said. And I don't know about you, but I think that's unbefrickinlievably awesome. ♣

This means, for example, that if $f\colon C_1 \to C_2$ is a rational map of curves, with C_1 smooth and C_2 projective, then f is actually a morphism! Since most singular curves are in prison for crimes against smoothness, this suggests that we should concentrate on projective curves to make our lives easier.

Let's take this opportunity to think a little more about morphisms of curves. First, notice that for any two curves C_1 and C_2, and any point $P \in C_2$, there is a constant map $f\colon C_1 \to C_2$ defined by $f(C_1) = P$. This is obviously very boring, and we will ghost it like a friend who refuses to buy the next round.

So now let's say we have a stylish *non-constant* rational map $f\colon C_1 \to C_2$. Here's a neat fact about it.

Theorem 1.3. *Let $f\colon C_1 \to C_2$ be a non-constant rational map of curves. For any point $P \in C_2$, the inverse image $f^{-1}(P) = \{Q \in C_1 \mid f(Q) = P\}$ is finite.*

Proof: The inverse image $f^{-1}(P)$ is an algebraic subset of C_1. (Its ideal is generated by the polynomials you get by plugging the coordinates of f into the variables in the defining polynomials for P.)

But C_1 has dimension 1 and is irreducible, so since f is non-constant, the irreducible components of $f^{-1}(P)$ must have dimension strictly smaller than 1. This means $f^{-1}(P)$ is a finite set of points. ♣

"Finite" is good. An actual number is better. But this requires a bit of added technology.

Definition 1.4. *Let $f\colon V \to W$ be a dominant rational map of algebraic varieties. The* **degree** *of f is the degree of the field extension $K(V)/f^*K(W)$, where $K(V)$ denotes the fraction field of V.*

The "dominant" is in there because otherwise you don't get a well defined pullback $f^*K(W)$.

It turns out that if V and W are curves, and if f is non-constant, then this degree is always finite. We won't take the time to prove this, but the key idea is that the Krull dimension of the coordinate ring is equal to the transcendence degree of the fraction field over $\mathbb{C}$. This means that $K(V)$ and $f^*K(W)$ have the same transcendence degree, and so the extension is algebraic – since it's also finitely generated, it must be finite.

If the previous paragraph stopped making sense after "We won't take the time to prove this", then don't worry. Just remember the fact:

Theorem 1.5. *The degree* $\deg f$ *of a dominant rational map of curves* $f\colon C_1 \to C_2$ *is finite.*

Anyway, we now have an actual number, namely, the degree of the map. Is it true that the inverse image of any point in C_2 has $\deg f$ points?

No.

Let $f\colon \mathbb{P}^1 \to \mathbb{P}^1$ be defined by $f(x : y) = [x^2 : y^2]$. Beautiful morphism. But the preimage of $[2 : 1]$ has two points – namely $[2 : 1]$ and $[2 : -1]$ – and the preimage of $[0 : 1]$ has only one point, namely $[0 : 1]$.

Darn.

(Think there are actually four points? Not so: $[-2 : 1]$ is the same point as $[2 : -1]$, because of scaling by -1.)

Well, I didn't introduce that degree for nothing. What's the degree of the f in our example?

Restricting to an affine piece will make calculations easier, and won't change the degree, which only depends on the fraction field.

And in the affine piece $y \neq 0$, the map f is $f(x) = x^2$. Meaning that the pullback morphism f^* embeds the function field $\mathbb{C}(x)$ of the destination $\mathbb{P}^1$ into the function field $\mathbb{C}(x)$ of the domain $\mathbb{P}^1$ by $f^*(p(x)) = p(x^2)$. So we have

$$[K(\mathbb{P}^1) : f^*K(\mathbb{P}^1)] = [\mathbb{C}(x) : \mathbb{C}(x^2)] = 2$$

and the map f has degree two.

And, y'know, the point $[x : y]$ usually has two preimages: $[\sqrt{x} : \sqrt{y}]$ and $[\sqrt{x} : -\sqrt{y}]$. It's only when x or y is zero that you get only one. So *most of the time*, our actual number is indeed the degree of f.

So can we prove this?

You bet your browser history we can. In fact, we can do better ... but we need (stop me if you've heard this before) more technology.

Definition 1.6. *Let $f\colon C_1 \to C_2$ be a dominant (i.e., non-constant) morphism of curves. Let $P \in C_1$ be any point of C_1, and let $Q = f(P) \in C_2$.*

Let t_Q be a uniformizer at Q, let t_P be a uniformizer at P, and write $f^*t_Q = ut_P^e$, where $u \in \mathcal{O}_P(C_1)$ is a unit.

The **ramification degree** *of f at P is the number e.*

Let's look at the previous example to illustrate this ramification degree thing. What's the ramification degree of f at $P = [2 : 1]$?

Well, $f(2 : 1) = [4 : 1]$ – let's set $Q = [4 : 1]$ – so let's specialize to the affine piece $y \neq 0$. Then $P = 2$ and $Q = 4$. A uniformizer at Q is $t_Q = x - 4$.

We have $f^*(t_Q) = f^*(x-4) = x^2 - 4 = (x-2)(x+2)$. Since $x+2$ does not vanish at 2, it's a unit in $\mathcal{O}_P(\mathbb{P}^1)$, and $t_P = x-2$ is a uniformizer at $P = 2$, so we have written

$$f^* t_Q = u t_P^1$$

and so $e = 1$.

Now let's try $P = [0:1]$. We have $Q = f(P) = [0:1]$, and so we again specialize to $y \neq 0$. A uniformizer at $Q = 0$ is $t_Q = x$, and we have $f^* t_Q = f^* x = x^2$. Since $t_P = x$ is a uniformizer at P, we have $e = 2$ at $P = 0$.

Interesting! Although the preimage of $[0:1]$ has only one point in it, its ramification degree is 2, so if you weight the preimage of $[0:1]$ by ramification degree, it *does* have two points in it!

I smell a theorem.

Theorem 1.7. *Let $f\colon C_1 \to C_2$ be a non-constant morphism of smooth projective curves. Then for every point $Q \in C_2$, we have*

$$\sum_{P \in f^{-1}(Q)} e_P = \deg f$$

where e_P denotes the ramification degree of f at P.

Proof: Tragically, the proof of this theorem requires either more algebra or more geometry than we have time to explain. So I'll pack all the mysterious algebra into a little black box, and prove the result from that.

The black box will come later. Don't worry – I'll tell you when it arrives.

Let $R \subset K(C_1)$ be the set of all rational functions on C_1 that are defined at all points $P_1, \ldots, P_r$ of $f^{-1}(Q)$. (Remember we already proved that $f^{-1}(Q)$ is finite.) Then we have

$$R = \mathcal{O}_{P_1}(C_1) \cap \ldots \cap \mathcal{O}_{P_r}(C_1)$$

Let $d = \deg f$. Then there is a basis $\phi_1, \ldots, \phi_d$ for $K(C_1)$ over $f^*K(C_2)$. Let T be a uniformizer for $f^*\mathcal{O}_Q(C_2)$. (Note that f^* is injective – it's a homomorphism of fields! – so $f^*\mathcal{O}_Q(C_2)$ is still a DVR.)

By judicious multiplication by powers of T to cancel out any poles that the ϕ_i might have at the points P_j, we can assume that all of the ϕ_i are defined at every point P_j. (Multiplying the ϕ_i by nonzero elements of $f^*K(C_2)$ – like T – doesn't change the fact that they're a basis of $K(C_1)$ over $f^*K(C_2)$.) In particular, this means that the fraction field of R is $K(C_1)$, because we can choose the ϕ_i to lie in R.

Black box time!

Theorem 1.8. *We can choose the ϕ_i so that every element r of R is a unique linear combination*

$$r = a_1\phi_1 + \ldots + a_d\phi_d$$

for $a_i \in f^*\mathcal{O}_Q(C_2)$.

Proof: Ok, well, we're not actually going to prove this theorem, because it's a black box. But I'll point out that the only thing that's nontrivial here is the ability to constrain the a_i to live in $f^*\mathcal{O}_Q(C_2)$. I mean, the ϕ_i are already a basis for $K(C_1)$ over $f^*K(C_2)$, so if you allow the a_i to be any old elements of $f^*K(C_2)$, we already know this theorem.

But it turns out to be surprisingly subtle to prove that you can always choose the a_i to lie in $f^*\mathcal{O}_Q(C_2)$.

Ok, fine. Here's a proof of the black box. But it involves some highfalutin algebra. Just sayin.

It turns out that R is the integral closure of $f^*\mathcal{O}_Q(C_2)$ in $K(C_1)$. To see this, notice first that R is integrally closed (it's the intersection of the integrally closed $\mathcal{O}_{P_i}(C_1)$), and it's integral over $f^*\mathcal{O}_Q(C_2)$ because the Galois action on elements of R respects fibres of f – that is, if $r \in R$, then any Galois conjugate of r will have zeros and poles (respectively) over the same points of C_2. In particular, none of the Galois conjugates of elements of r have poles over Q, and so the coefficients of the monic minimal polynomial for r over $f^*K(C_2)$ all lie in $f^*\mathcal{O}_Q(C_2)$.

So R is the integral closure of $f^*\mathcal{O}_Q(C_2)$ in $K(C_1)$. That means it's a finitely generated, torsion free module over $f^*\mathcal{O}_Q(C_2)$. Since $f^*\mathcal{O}_Q(C_2)$ is a DVR, it's also a PID, so R is a *free* module over $f^*\mathcal{O}_Q(C_2)$. Which is exactly what the Black Box Theorem is saying.

And just to be clear: *I did not expect you to understand any of the proof beyond the word "sayin"*. ♣

So, as an additive group, R is isomorphic to $(f^*\mathcal{O}_Q(C_2))^d$. All we gotta do now is show that $d = \sum_{P \in f^{-1}(Q)} e_P$.

Remember that by the Black Box, we can write any element of R uniquely in the following way:

$$r = a_1\phi_1 + \ldots + a_d\phi_d$$

for $a_i \in f^*\mathcal{O}_Q(C_2)$. Better yet, for any element $x \in R$, it's also true that r is in the ideal xR if and only if all the coefficients a_i

lie in TR.

To see this, first notice that if all the a_i lie in xR, then obviously

$$r = a_1\phi_1 + \ldots + a_d\phi_d \in xR$$

And conversely, if

$$r = a_1\phi_1 + \ldots + a_d\phi_d = ax \in xR$$

for some $a \in R$, then

$$a = \frac{a_1}{x}\phi_1 + \ldots + \frac{a_d}{x}\phi_d \in R$$

and so, by the Black Box, we have $a_i/x \in R$ for all i, or equivalently, $a_i \in xR$ for all i. Awesome.

Now we can apply this when $x = T$ is a uniformizer for $f^*\mathcal{O}_Q(C_2)$. What the previous paragraph means is that the ϕ_i are linearly independent mod TR. And R/TR is a vector space over $f^*\mathcal{O}_Q(C_2) \cong \mathbb{C}$. Let's compute its dimension. (Hint: its dimension is $\sum_{P \in f^{-1}(Q)} e_P$.)

Well, by the definitions of the ϕ_i and e_i, we have $T = u\psi_1^{e_1}\ldots\psi_r^{e_r}$ for some unit u of R. And ψ_i is coprime to ψ_j if $i \neq j$, because they both generate maximal ideals of R. (The quotient is isomorphic to $\mathbb{C}$ – the quotient map is evaluation at the point P_i. Or P_j, depending on which ideal you're talking about.)

Anyway, the coprimality means that we can apply the Chinese Remainder Theorem:

$$R/TR \cong \prod_{i=1}^{r} R/(\psi_i^{e_i})R$$

so that the dimension of R/TR as a complex vector space is just the sum of the dimensions of the $R/(\psi_i^{e_i})R$.

So it remains only to show that

$$\dim R/(\psi_i^{e_i})R = e_i$$

Remember this theorem?

Theorem 1.9. *Let D be a domain, D_M the localization of D at a maximal ideal M. For any positive integer n, the natural inclusion $M \hookrightarrow \mathcal{M}$ induces an isomorphism*

$$M^{n-1}/M^n \cong \mathcal{M}^{n-1}/\mathcal{M}^n$$

where $\mathcal{M}$ denotes the ideal of D_M generated by M. In particular, for $n = 1$ and $n = 2$, we have

$$D/M \cong D_M/D_M/\mathcal{M}$$

and

$$M/M^2 \cong \mathcal{M}/\mathcal{M}^2$$

Well, $\mathcal{O}_{P_i}(C_1)$ is the localization of R at the maximal ideal $\psi_i R$. (Check it out if you don't believe me. If you can't figure out the short proof, ask me and I'll show you.)

Which means that the complex dimension of $R/\psi_i R$ is equal to the complex dimension of $\mathcal{O}_{P_i}(C_1)/\psi_i\mathcal{O}_{P_i}(C_1)$, which is to say, dimension one. (The quotient map is, again, plugging in P_i.)

It's a simple induction to show the following equality of dimensions:

$$\dim \mathcal{O}_{P_i}(C_1)/\psi_i^e\mathcal{O}_{P_i}(C_1) = \dim \mathcal{O}_{P_i}(C_1)/\psi_i\mathcal{O}_{P_i}(C_1) + \dots$$
$$\dots + \dim \psi^{e-1}\mathcal{O}_{P_i}(C_1)/\psi_i^e\mathcal{O}_{P_i}(C_1)$$

Each of the dimensions on the right side is one. So the dimension on the left side is e.

Which means that we're done! ♣