Contents

0 Number Theory in Function Fields

These are lecture notes for the graduate course Math 7360: Number Theory in Function Fields, taught at Northeastern in Fall 2025.

0.1 (Sep 3) Overview + Fermat's Last Theorem for Polynomials

- The goal of this course is to elucidate some of the many analogies between number theory in number fields and number theory in function fields.
    - Some things from classical number theory: primes, factorizations, congruences and modular arithmetic, Fermat's and Euler's theorems, the prime number theorem, quadratic reciprocity (and higher reciprocity), Dirichlet's theorem on primes in arithmetic progressions, zeta functions.
    - Some things from the more modern take on algebraic and analytic number theory: algebraic number fields and their rings of integers, Galois theory and its interplay with number fields, discriminants, class groups, Dirichlet's unit theorem, cyclotomic fields, ramification, L-functions, the Riemann hypothesis.
    - Our goal is to do as much of these things as possible in the context of function fields, where many of the results are more approachable, because the function-field setting has a major kit of additional tools (namely, algebraic geometry).
    - Though do note: number theory in function fields is a beautiful subject in its own right, and not just because it has so many similarities to algebraic number theory.
    - We will illustrate how things can become simpler by proving Fermat's Last Theorem, which is quite notoriously difficult over $\mathbb{Z}$, for polynomials using only elementary techniques.

- To start, let $q = p^f$ be a prime power, and let $\mathbb{F}_q$ be the finite field with q elements. The story begins with the polynomial ring $A = \mathbb{F}_q[t]$.
  - We have the degree map on A: explicitly, for coefficients $a_i \in \mathbb{F}_q$ and an element $f = a_0 + a_1 t + \cdots + a_n t^n$ with $a_n \neq 0$, we define $\deg(f) = n$ and $\operatorname{sgn}(f) = a_n$. (We also set $\deg(0) = -\infty$ and $\operatorname{sgn}(0) = 0$.)
  - Exercises (trivial): $\deg(fg) = \deg(f) + \deg(g)$, $\operatorname{sgn}(fg) = \operatorname{sgn}(f)\operatorname{sgn}(g)$, and $\deg(f+g) \leq \max(\deg f, \deg g)$ with equality whenever $\deg f \neq \deg g$.
  - The polynomials with sign 1 (i.e., monic polynomials) behave analogously to the integers with positive sign (i.e., the positive integers).
  - We also note that the degree properties easily give a characterization of the units of A: they are the nonzero constant polynomials.
- Our first basic result is the standard division-with-remainder algorithm for polynomials, which we record over arbitrary fields for no extra cost:
- Exercise (Polynomial Division): If F is any field, then for any $f, g \in F[t]$ with $g \neq 0$, there exist unique $q, r \in F[t]$ such that $f = qg + r$ and $\deg r < \deg g$.
  - The idea is simply to prove that the usual long-division algorithm works by induction on the degree of g.
  - As a consequence, $F[t]$ is a Euclidean domain, meaning that it is also a principal ideal domain (all ideals are principal) and a unique factorization domain (every element can be factored uniquely into a product of irreducibles up to reordering and unit factors).
- As it turns out, unique factorization is essentially all we need to prove Fermat's Last Theorem for polynomials.
  - We would like to show that the equation $f^n + g^n = h^n$ has no nontrivial solutions in polynomials f, g, h. Aside from the case $n = 4$, it is enough to treat the situation where n is a prime.
  - But we do need to be a little bit careful to write down exactly what the trivial solutions look like, beyond the obvious ones where one of f, g, h is zero.
  - For example, if f, g, h are all constants, we can certainly have lots of solutions to $f^n + g^n = h^n$, depending on the field and on n (e.g., $1^5 + 1^5 = 2^5$ inside $\mathbb{F}_3$).
  - We need to avoid the situation where n is divisible by $p = \operatorname{char}(\mathbb{F}_q)$, since $f^p + g^p = (f+g)^p$ for any polynomials $f, g \in \mathbb{F}_q[t]$.
  - Also, since the equation is homogeneous, we can scale solutions to get new solutions.
  - To avoid all of these situations, we can consider only the case where f, g, h are relatively prime (since if they are not, then any common divisor of two of them also divides the third, so we could cancel it) and where the exponent n is not divisible by the characteristic p.
- Theorem (FLT for Polynomials): Suppose that $f, g, h \in F[t]$ are pairwise relatively prime and that $p \geq 3$ is prime with $p \neq \operatorname{char}(F)$. Then the only solutions to $f^p + g^p = h^p$ are when f, g, h are all constants.
  - We will remark that $p \geq 3$ is needed, since the usual parametrization of Pythagorean triples also works for polynomials: if we take $f = a^2 - b^2$, $g = 2ab$, $h = a^2 + b^2$ for any polynomials $a, b \in F[t]$, then $f^2 + g^2 = h^2$.
- We will give two different proofs: the first uses a classical-style infinite descent argument, while the second uses a more function-field type of argument.
  - Proof 1: Without loss of generality, we may assume that F is algebraically closed, since any solution to $f^p + g^p = h^p$ over F is still a solution over the algebraic closure $\overline{F}$.
  - We show the result by inducting on $d = \deg f + \deg g$. The base case $d = 0$ is trivial, since there is nothing to prove. So now suppose we have a solution with $d > 0$.
  - By the assumption that $p \neq \operatorname{char}(F)$, there are p distinct pth roots of unity in F: say, $1, \zeta_p, \zeta_p^2, \ldots, \zeta_p^{p-1}$, and we can factor $f^p + g^p = (f+g)(f+\zeta_p g)(f+\zeta_p^2 g)\cdots(f+\zeta_p^{p-1} g)$.

- Next, note that all of the terms $f+\zeta_p^i g$ are relatively prime: if e divides both $f+\zeta_p^i g$ and $f+\zeta_p^j g$, then e also divides the difference $(\zeta_p^i-\zeta_p^j)g$ hence divides g, hence also divides $(f+\zeta_p^i g)-\zeta_p^i g = f$, but f and g are relatively prime by assumption.
- Then by unique factorization inside $F[t]$, since all of the terms in the product $(f+g)(f+\zeta_p g)(f+\zeta_p^2 g)\cdots(f+\zeta_p^{p-1}g)$ are relatively prime and their product is a pth power (namely, h^p), each term must be a pth power up to a unit factor. But since F is algebraically closed, everything in F has a pth root in F, so the unit factor is also a pth power.
- Thus, in particular, we see that $f+g=a^p$, $f+\zeta_p g = b^p$, and $f+\zeta_p^2 g = c^p$ are all pth powers.
- Using basic linear algebra to eliminate f and g yields the relation $-\zeta_p a^p + (1+\zeta_p)b^p = c^p$, so if we set $a' = (-\zeta_p)^{1/p}a$, $b' = (1+\zeta_p)^{1/p}b$, and $c'=c$, then we have $(a')^p+(b')^p=(c')^p$.
- Note that a', b' cannot both be constant, since then f, g would have been constant. But we also have $\deg(a')+\deg(b') = \deg(f+g)/p + \deg(f+\zeta_p g)/p \le 2\max(\deg f, \deg g)/p < d$, so we have constructed a solution with smaller positive degree, but this contradicts the induction hypothesis. Therefore, there are no nonconstant solutions.
- Exercise: For any field F of characteristic p, we have exhibited nontrivial polynomial solutions to $f^p+g^p = h^p$ in $F[t]$. Where and why in the proof of FLT above does the argument break down when $\mathrm{char}(F)=p$?

- Before giving the second proof, we need a few preliminary results.
  - First, if f has prime factorization $f=\prod_i p_i^{a_i}$, define $\mathrm{rad}(f)=\prod_i p_i$, the product of the monic irreducible polynomials dividing f.

- Lemma: We have $\deg\gcd(f,f') \ge \deg f - \deg \mathrm{rad} f$, where f' is the derivative of f.
  - Proof: Suppose $f=p^a q$ where p is irreducible and doesn't divide q. Then $f' = ap^{a-1}p'q + p^a q' = p^{a-1}(ap'q+pq')$ is divisible by p^{a-1}. Therefore, $\gcd(f,f')$ is divisible by p^{a-1}.
  - Taking the product over all primes dividing f shows that $\prod_i p_i^{a_i-1}$ divides $\gcd(f,f')$, so $\gcd(f,f')\cdot\mathrm{rad}(f)$ is divisible by $\prod_i p_i^{a_i-1}\prod_i p_i = \prod_i p_i^{a_i} = f$, so taking degrees yields the inequality.
  - Exercise: Determine when equality holds, namely when $\deg\gcd(f,f') = \deg f - \deg\mathrm{rad} f$.

- Next, we show a result due independently to Mason and Stothers:

- Proposition (Mason-Stothers): Suppose that $f,g,h \in F[t]$ are nonconstant, relatively prime, that $f+g=h$, and that not all of f',g',h' are zero. Then $\max(\deg f, \deg g, \deg h) \le \deg\mathrm{rad}(fgh)-1$.
  - Proof: If $f+g=h$ then $f'+g'=h'$, and then $fg'-f'g = (f+g)g' - (f'+g')g = hg'-h'g$.
  - Note also that $fg'-f'g$ is nonzero: if $fg'=f'g$ then f must divide $f'g$ hence that f must divide f' since f,g are relatively prime.
  - Exercise: Suppose $f\in F[t]$. Show that f divides its derivative f' if and only if $f'=0$.
  - By the exercise we see then that $f'=0$. But now by the same argument we would also have $g'=0$ and $h'=0$, contradicting the assumption that not all of f',g',h' are zero.
  - Now let $d_f=\gcd(f,f')$, $d_g=\gcd(g,g')$, $d_h=\gcd(h,h')$. Then d_f, d_g, d_h all divide $fg'-f'g = hg'-h'g$, and they are all relatively prime since they are divisors of the relatively prime polynomials f,g,h.
  - This means $d_f d_g d_h$ divides $fg'-f'g$, so taking degrees yields $\deg(d_f d_g d_h) \le \deg(fg'-f'g) \le \deg(f)+\deg(g)-1$.
  - By the Lemma, we have $\deg(d_f)\ge\deg(f)-\deg\mathrm{rad}f$, $\deg(d_g)\ge\deg(g)-\deg\mathrm{rad}g$, $\deg(d_h)\ge\deg(h)-\deg\mathrm{rad}h$, so summing yields $\deg(f)+\deg(h)+\deg(h)-\deg\mathrm{rad}(fgh) \le \deg(d_f d_g d_h) \le \deg(f)+\deg(g)-1$, and therefore $\deg(h)\le\deg\mathrm{rad}(fgh)-1$.
  - By rearranging we obtain the same bounds on $\deg(f)$ and $\deg(g)$, and so we are done.

- At last, we can finish the second proof of Fermat's Last Theorem for polynomials:
  - Proof 2: Suppose $f^p+g^p=h^p$. By the assumption on the characteristic, we have $(f^p)'$, $(g^p)'$, $(h^p)'$ are not all zero.

- Then by Mason-Stothers, we see $\max(\deg f^p, \deg g^p, \deg h^p) \leq \deg \operatorname{rad}(f^p g^p h^p) - 1$, which is equivalent to $p \cdot \max(\deg f, \deg g, \deg h) \leq \deg \operatorname{rad}(fgh) - 1 \leq \deg f + \deg g + \deg h - 1$ since the radical ignores powers.
- Now apply the simple observation that $\max(a, b, c) \geq (a + b + c)/3$ and set $d = \deg f + \deg g + \deg h$ to see that $p \cdot d/3 \leq d - 1$, which is impossible, since $d \leq p \cdot d/3$ by the hypothesis that $p \geq 3$.

0.2 (Sep 8) Quotients of $\mathbb{F}_q[t]$

- We now return to study the structure of quotient rings of $A = \mathbb{F}_q[t]$, which (re-posed) is simply studying modular arithmetic in this ring.
  - In particular, we will recover almost identical versions of Fermat's little theorem, Euler's theorem, and Wilson's theorem.
  - We will also take some time to look at the structure of the unit group of A/gA, which turns out to be a bit more complicated to write down than the unit group of $\mathbb{Z}/m\mathbb{Z}$.
- As noted last lecture, A is a Euclidean domain, so it is a PID and also a UFD. Since every ideal is principal, if we want to understand the structure of the quotient rings of A, we only have the quotients of the form A/gA to consider.
  - We can also assume g is monic by replacing it with its unique monic associate, which does not change the quotient ring A/gA.
- Using the division algorithm, we can write down the residue classes in A/gA, and in particular compute its cardinality, quite easily:
- Proposition: Let $g \in \mathbb{F}_q[t] = A$ be nonzero. Then the residue classes in A/gA are uniquely represented by the polynomials of degree less than $\deg(g)$. In particular, $\#(A/gA) = q^{\deg g}$.
  - Proof: If $f \in \mathbb{F}_q[t]$ is any polynomial, then by the division algorithm we can write $f = qg + r$, and so inside A/gA we see $\overline{f} = \overline{r}$. So the possible remainders give a complete set of residue class representatives – but by the uniqueness of the quotient and remainder, no two remainders are equivalent mod g, so in fact they give all of the residue classes exactly once.
  - For the counting, if $\deg(g) = n$, then the remainders are of the form $c_0 + c_1 t + \cdots + c_{n-1} t^{n-1}$ with $c_i \in \mathbb{F}_q$. Since there are n coefficients each of which has q possible values, there are $q^n = q^{\deg g}$ possible ways to select a remainder.
- The size of the quotient ring gives a convenient way of measuring the "size" of a polynomial that behaves pleasantly under multiplication:
- Definition: For $g \in \mathbb{F}_q[t]$, we define $|g|$, the norm of g, to be $q^{\deg g}$. By the calculation above, $|g| = \#(A/gA)$ when $g \neq 0$.
  - Exercise: Show $|fg| = |f| \cdot |g|$ and $|f + g| \leq \max(|f|, |g|)$ with equality whenever $|f| \neq |g|$.
- Our next goal is to understand the units of A/gA, since this is the context in which to pose Fermat's and Euler's theorems.
  - Regardless of the polynomial g, the units of A/gA will contain an isomorphic copy of the constant polynomials (i.e., the units of A), which is the multiplicative group $\mathbb{F}_q^*$.
  - As is well-known, the multiplicative group of a finite field is cyclic. We record a few proofs of this fact, for completeness:
- Proposition (Multiplicative Group of $\mathbb{F}_q$): If G is a finite multiplicative subgroup of a field F, then G is cyclic.
  - All known proofs of this fact are essentially nonconstructive, to varying degrees: there does not seem to be a nice algorithm for writing down a multiplicative generator of a finite field that is appreciably better than a brute-force search.

- Proof 1: Let G be a finite multiplicative subgroup of F. By the fundamental theorem of finite(ly generated) abelian groups, G is isomorphic to a direct product of cyclic groups.
- Let m be the lcm of the orders of these cyclic groups: then $x^m = 1$ for all $x \in G$. Since $F[t]$ has unique factorization, the polynomial $t^m - 1 \in F[t]$ has at most m roots in F, so $\#G \leq m$. On the other hand, by Lagrange's theorem, the order of every element in G divides $\#G$, so m divides $\#G$. We must therefore have $m = \#G$.
- But since $\#G$ is equal to the product of the orders of the cyclic groups, we see that the product of these orders equals their lcm, so the orders are all relatively prime. This means G is cyclic, as claimed.
- Proof 2: Let M be the maximal order among all elements in G: we claim that the order of every element in G divides M. To see this, suppose g has order M, and let h be any other element of order k. If k does not divide M, then there is some prime q which occurs to a higher power q^f in the factorization of k than the corresponding power q^e dividing M.
- By properties of orders, the element g^{q^f} has order M/q^f, and the element h^{k/q^e} has order q^e. Since these two orders are relatively prime and $gh = hg$ (since these are elements in a field), we see that the element $g^{q^f} \cdot h^{k/q^e}$ has order $M \cdot q^{f-e}$. This is a contradiction because this element's order is larger than M. Thus, k divides M as claimed.
- For the second claim, any element of order M generates a subgroup of G having M elements, so $M \leq \#G$.
- Furthermore, by the above, we know that all elements in G have order dividing M, so the polynomial $t^M - 1$ has $\#G$ roots in $F[t]$. By unique factorization, this requires $M \geq \#G$, and so we have $M = \#G$. Now select any element of order M: it generates G.
- Proof 3: Observe by Lagrange's theorem that $t^{\#G} - 1$ factors as the product $\prod_{d|\#G} \Phi_d(t)$, where $\Phi_d(t) = \prod_{\text{order(g)=d}}(t - g)$ is the dth cyclotomic polynomial.
- By an inductive argument, or by observing invariance under the Galois action, all of the polynomials $\Phi_d(t)$ have coefficients in $F[t]$.
- By induction on d using the fact that $t^d - 1$ has at most (hence exactly) d roots in F and in G, one has that $\deg(\Phi_d) = \varphi(d)$. In particular, $\deg(\Phi_{\#G}) = \varphi(\#G) > 0$, so there is an element of order $\#G$ in G.

- Now we tackle the question of the units of A/gA.
  - We can simplify the problem first: if we factor $g = p_1^{a_1} \cdots p_d^{a_d}$ where the p_i are distinct monic irreducible polynomials, then all of the ideals $(p_i^{a_i})$ are pairwise comaximal, so by the Chinese remainder theorem, we see $A/gA \cong (A/p_1^{a_1}A) \times (A/p_2^{a_2}A) \times \cdots \times (A/p_d^{a_d}A)$.
  - Taking units on both sides then gives $(A/gA)^* \cong (A/p_1^{a_1}A)^* \times (A/p_2^{a_2}A)^* \times \cdots \times (A/p_d^{a_d}A)^*$. So it is enough to study the structure of the ring A/p^aA where p is irreducible.

- Proposition (Structure of A/p^aA): For $A = \mathbb{F}_q[t]$ where $\text{char}(\mathbb{F}_q) = \tilde{p}$, and $p \in A$ is a monic irreducible polynomial, we have the following:

  1. The cardinality of $(A/p^aA)^*$ is $\#(A/p^aA)^* = |p|^{a-1}(|p| - 1) = |p^a|(1 - 1/|p|)$.
     - Exercise: Show that a commutative ring R with 1 has a unique maximal ideal M if and only if the set of nonunits in R forms an ideal, which is then a unique maximal ideal M. A ring with this property is called a local ring.
     - Proof: The ring A/p^aA has a unique maximal ideal, namely pA/p^aA, and is therefore a local ring, because the quotient $(A/p^aA)/(pA/p^aA) \cong A/pA$ is a field by the third isomorphism theorem.
     - By the exercise above, evvery element not in the maximal ideal is a unit, and the cardinality of the maximal ideal is $1/|p|$ times the cardinality of the entire ring (since the elements in the ideal are just the multiples of p). The formula follows.
  2. $(A/p^aA)^* \cong$ [cyclic group of order $|p| - 1$] $\times$ [an abelian $\tilde{p}$-group].
     - Proof: The reduction-mod-p map is a surjective group homomorphism from $(A/p^aA)^* \to (A/pA)^*$, and the latter is the multiplicative group of the field A/pA hence is cyclic of order $|p| - 1$.
     - Pulling back a generator yields that $(A/p^aA)^*$ contains a cyclic subgroup of order $|p| - 1$. By the cardinality calculation in (1), the remaining piece has order $|p|^{a-1}$ and is therefore a $\tilde{p}$-group (and it is clearly abelian).

- Remark: The direct product decomposition writes each element modulo p^a as [its residue modulo p] times [an element congruent to 1 modulo p].

3. The $\tilde{p}$-part of $(A/p^aA)^*$ has exponent at most $\tilde{p}^s$ where $\tilde{p}^s \geq a$.
   - Proof: By the above, the elements in the $\tilde{p}$-part are of the form $1+bp$ for some $b \in \mathbb{F}_q[t]$.
   - Since we are in characteristic $\tilde{p}$, we then have $(1+bp)^{\tilde{p}^s} = 1+(bp)^{\tilde{p}^s}$, and since $p^{\tilde{p}^s}$ is divisible by p^a by assumption, we see $(1+bp)^{\tilde{p}^s} \equiv 1 \pmod{p^a}$, which is to say, the element $1+bp$ modulo p^a has order dividing $\tilde{p}^s$ (as required).
4. As $a \to \infty$, the number of cyclic factors in the $\tilde{p}$-part of $(A/p^aA)^*$ goes to infinity.
   - The point here is that we get a different kind of behavior than over $\mathbb{Z}$: over $\mathbb{Z}$, we see that $(\mathbb{Z}/p^a\mathbb{Z}) \cong \begin{cases} \mathbb{Z}/(p^a - p^{a-1})\mathbb{Z} & \text{for odd primes } p \\ (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2^{a-3}\mathbb{Z}) & \text{for } p = 2 \end{cases}$, and so even for large prime powers, the quotient is either cyclic or basically cyclic.
   - For polynomials, we end up getting a large number of cyclic factors when we take a large power, regardless of the prime.
   - Proof: Since the exponent of the $\tilde{p}$-part is at most $\tilde{p}^s$, if we have a total of j cyclic factors then the order of the group is at most $\tilde{p}^{sj}$. So we need $\tilde{p}^{sj} \geq |p|^{a-1} = q^{\deg(p)\cdot(a-1)} = \tilde{p}^{f\cdot\deg(p)\cdot(a-1)}$ and so $j \geq f \cdot \deg(p) \cdot (a-1)/s$.
   - Since $s \sim \log_p a$, we see that for a fixed field $\mathbb{F}_q$ (i.e., fixed f) and fixed prime p (i.e., fixed $\deg p$), we have $j \sim C(a-1)/\log_p a \to \infty$ as $a \to \infty$.

- Now that we have established some basic things about the unit group of A/p^aA, we can establish the analogues of Fermat's little theorem, Euler's theorem, and Wilson's theorem.
  - First, we need the analogue of the Euler phi-function. We define $\Phi(f) = \#(A/fA)^*$ to be the number of polynomials of degree less than $\deg f$ that are relatively prime to f.
  - By our calculations with the unit group earlier, we have the usual formula $\Phi(f) = |f| \prod_{p|f \text{prime}}(1 - 1/|p|)$, which is the analogue of $\varphi(n) = n\prod_{p|n \text{ prime}}(1-1/p)$ for the phi-function over $\mathbb{Z}$.

- Proposition ("Euler"): If $f \in \mathbb{F}_q[t]$ is nonzero and g is relatively prime to f, then $g^{\Phi(f)} \equiv 1 \pmod{f}$.
  - Proof 1: Apply Lagrange's theorem to $\overline{g}$ in $(A/fA)^*$.
  - Proof 2: Multiplication by $\overline{g}$ is a bijection on the cosets in $(A/fA)^*$. Thus, $\prod_{u\in(A/fA)^*} u = \prod_{u\in(A/fA)^*}(ug) = g^{\Phi(f)}\prod_{u\in(A/fA)^*} u$ inside $(A/fA)^*$, and cancelling the unit factor $\prod_{u\in(A/fA)^*} u$ yields $g^{\Phi(f)} = 1$ inside $(A/fA)^*$.

- Proposition ("Fermat"): If $p \in \mathbb{F}_q[t]$ is irreducible, then $a^{|p|} \equiv a \pmod{p}$ for any $a \in \mathbb{F}_q[t]$.
  - Proof: If $p|a$ the result is trivial. Otherwise, a is a unit modulo p and the result follows from Euler above.

- We can use the analogue of Fermat's theorem to prove an analogue of Wilson's theorem:

- Proposition (Factoring, 1): If $p \in \mathbb{F}_q[t]$ is irreducible of degree d, then $x^{|p|} - x \equiv \prod_{\deg f < d}(x - f) \bmod p$.
  - Proof: As we have noted, in A/p the polynomials of degree $< d$ represent all of the residue classes modulo p.
  - By Fermat, each of these polynomials is a root of $x^{|p|} - x$. But by unique factorization, this polynomial has at most $|p|$ distinct roots, and we have just exhibited $|p|$ roots, so these are all of the roots, and the factorization follows.

- Corollary ("Wilson"): If $p \in \mathbb{F}_q[t]$ is irreducible of degree d, then $\prod_{\deg f<d, f\neq 0} f \equiv -1 \pmod{p}$.
  - Proof 1: Dividing the result above by x yields $x^{|p|-1} - 1 \equiv \prod_{\deg f<d, f\neq 0}(x - f) \bmod p$.
  - Now set $x = 0$: if the characteristic is odd, then the number of minus signs on the RHS is even and the result follows, while if the characteristic is even, then $1 = -1$ so the result still follows.

- Proof 2: If $\overline{f}$ does not have order 2 in A/pA, then $\overline{f} \neq \overline{f}^{-1}$ and so we can pair up and discard $(\overline{f}, \overline{f}^{-1})$ without affecting the product.
- When we have done this for all possible pairs, the only elements left are the elements of order dividing 2 (i.e., the solutions to $x^2 = 1$), which are $x = \pm 1$. In characteristic not 2, the product is -1, while in characteristic 2, the product is $1 = -1$.
- Exercise: Generalize proof 2 of Wilson's theorem to show that if G is a finite abelian group, then the product of all elements in g is the unique element in G of order 2 (if there is one), or is otherwise the identity.

- We also record a useful result about roots of unity:
- Proposition (Roots of Unity): If $p \in \mathbb{F}_q[t] = A$ is irreducible and d divides $|p| - 1$, then there are d dth roots of unity in A/pA; equivalently, $x^d \equiv 1 \pmod{p}$ has exactly d solutions.
  - Exercise: For positive integers a, b, show $\gcd(x^a - 1, x^b - 1) = x^{\gcd(a,b)} - 1$ in $F[x]$.
  - Proof: As shown above, $x^{|p|-1} - 1$ splits completely mod p . By the exercise, $x^d - 1$ divides $x^{|p|-1} - 1$ when d divides $|p| - 1$, and so $x^d - 1$ also splits completely, which is to say, it has d roots mod p.
  - Exercise: Prove the converse: if there are d dth roots of unity in A/pA, then d divides $|p| - 1$.

0.3 (Sep 10) Prime-Counting and The Zeta Function

- Now that we have established most of the classical results for modular arithmetic, we move to our next item: counting primes.
  - We will do things in a more ad hoc manner first, and then give a more general approach using zeta functions that will allow us to go further.

- Our first step is to write down a generalization of the fact we used to establish Wilson's theorem above:
- Theorem (Factoring, II): For a positive integer m, the polynomial $t^{q^m} - t$ factors in $\mathbb{F}_q[x]$ as the product of all monic irreducible polynomials of degree dividing m.
  - Proof 1 ("Elementary"): We will show that $t^{q^m} - t$ has no repeated factors, that each of the claimed polynomials does divide it, and that no other polynomials divide it.
  - Exercise: A polynomial in $F[t]$ is separable (i.e., has no repeated factors) if and only if it is relatively prime to its derivative.
  - Since $(t^{q^m} - t)' = q^m t^{q^m - 1} - 1 = -1$ in characteristic p, the polynomial is relatively prime to its derivative, so it has no repeated factors by the exercise.
  - Exercise: For positive integers q, a, b, show that $\gcd(q^a - 1, q^b - 1) = q^{\gcd(a,b)} - 1$ in $\mathbb{Z}$. (This is almost identical to the polynomial version mentioned earlier.)
  - Next, suppose p is irreducible of degree dividing m. If $p = t$ the result is trivial, and otherwise, in A/pA we have $t^{q^m - 1} \equiv 1$ mod p because $q^m - 1$ is a multiple of $|p| - 1 = q^{\deg p} - 1$ by the exercise above along with Euler's theorem. This means $t^{q^m - 1} - 1$ is divisible by p as required.
  - Finally, suppose p is irreducible of degree not dividing m. Then in A/pA we have $t^{q^m - 1} \equiv t^{q^{\gcd(m, \deg p)}} \neq 1$ mod p by the exercise above along with Euler's theorem and the fact that $q^{\gcd(m, \deg p)} < q^{\deg p}$. This means $t^{q^m - 1} - 1$ is not divisible by p as required.
  - We have shown that $t^{q^m} - t$ has no repeated factors, that each of the claimed polynomials does divide it, and that no other polynomials divide it. Since the polynomial is monic, its factorization must therefore be as claimed.
  - Proof 2 ("Galois"): By basic Galois theory, $\mathrm{Gal}(\mathbb{F}_{q^m}/\mathbb{F}_q)$ is a cyclic group of order m generated by the Frobenius map $x \mapsto x^q$[1].

[1] This follows by noting that $\mathbb{F}_{q^m}$ is the splitting field of $x^{q^m} - x$ over $\mathbb{F}_q$ and since this polynomial is separable as noted in proof 1, the order of the Galois group is m. The Frobenius map is an injective field map from $\mathbb{F}_{q^m}$ to itself, hence an automorphism by finiteness, and its order is clearly at least m (since $x^{q^d} - x$ has at most q^d solutions) and at most m (by Lagrange).

- By the Galois correspondence, the intermediate fields of $\mathbb{F}_{q^m}/\mathbb{F}_q$ are $\mathbb{F}_{q^d}$ for $d|m$. Therefore, p is irreducible of degree dividing $d \iff \mathbb{F}_q[t]/(p)$ is (isomorphic to) an intermediate field of $\mathbb{F}_{q^m}/\mathbb{F}_q \iff p$ divides $x^{q^m} - x$.
- Since $x^{q^m} - x$ is separable, its factorization must therefore be as claimed.

- Corollary: If a_d is the number of irreducible monic polynomials in $A = \mathbb{F}_q[t]$ of degree d, then $\sum_{d|n} da_d = q^n$.
  - Proof: Count degrees in the theorem above.
- We can use this recurrence to write down an exact formula for a_d using Mobius inversion.
- Definition: The Mobius μ-function is defined as $\mu(n) = \begin{cases} 0 & \text{if } n\text{is not squarefree} \\ (-1)^r & \text{if } n\text{is the product of } r\text{distinct primes} \end{cases}$. Note $\mu(1) = 1$.
  - Exercise: Show that $\sum_{d|n} \mu(d) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$.
- Proposition (Mobius Inversion): If f, n are integer functions such that $g(n) = \sum_{d|n} f(d)$, then $f(n) = \sum_{d|n} \mu(d) g(n/d)$.
  - Proof: Induct on n. The base case $n = 1$ is trivial.
  - For the inductive step, we have $\sum_{d|n} \mu(d) g(n/d) = \sum_{d|n} \mu(d) \cdot \sum_{d'|n/d} f(d') = \sum_{dd'|n} \mu(d) f(d') = \sum_{d'|n} f(d') \sum_{d|(n/d')} \mu(d) = f(n)$ because the last inner sum is zero except for when $n/d' = 1$.
- By using Mobius inversion on the sequence $\{da_d\}$, we can write down formulas for the number of monic irreducible polynomials of degree d.
- Proposition (Prime Counting): If a_n is the number of monic irreducible polynomials in $\mathbb{F}_q[t]$ of degree n, then $a_n = \dfrac{1}{n} \sum_{d|n} \mu(d) q^{n/d}$.
  - The first few values are $a_1 = q$, $a_2 = \frac{1}{2}(q^2 - q)$, $a_3 = \frac{1}{3}(q^3 - q)$, $a_4 = \frac{1}{4}(q^4 - q^2)$, $a_5 = \frac{1}{5}(q^5 - q)$, $a_6 = \frac{1}{6}(q^6 - q^3 - q^2 + q)$, ...
  - Proof: Immediate from applying Mobius inversion to the sequence $\{na_n\}$.
- We can also do some basic asymptotic analysis using the formula above.
  - The main term is $\dfrac{1}{n} q^n$, and then the next biggest possible term is $\dfrac{1}{n} q^{n/2}$, so we see that $a_n = \dfrac{1}{n} q^n + O(q^{n/2}/n)$.
  - If we write $X = q^n$ (which is the total number of monic polynomials of degree n), we see that the number of "primes" in A of "size" $\sim X$ is $a_n = \dfrac{X}{\log_q X} + O(\dfrac{\sqrt{X}}{\log_q X})$.
  - This is quite in the spirit of the prime number theorem over $\mathbb{Z}$, which says that the number of primes $\leq X$ is $\Pi(X) = \dfrac{X}{\log X} + O(\dfrac{X}{(\log X)^2})$. If we replace $X/\log X$ with the logarithmic integral $\mathrm{li}(x) = \int_2^x \dfrac{dt}{\log t}$, then as shown by von Koch, the Riemann hypothesis is equivalent to the error estimate $\Pi(X) = \mathrm{li}(x) + O(\sqrt{X} \log x)$.
  - Qualitatively, then, we have already obtained a prime-counting result that is closely analogous to the best possible one predicted by the Riemann hypothesis.
- Up until this point, our approach has been purely algebraic. However, by introducing analytic methods, we can give even easier solutions to these (and other) counting problems. The necessary object of study is the zeta function, which we now define:
- Definition: For $A = \mathbb{F}_q[t]$, the zeta function of A is $\zeta_A(s) = \sum_{f \in A \text{ monic}} \dfrac{1}{|f|^s}$ for $s \in \mathbb{C}$.

- Compare to the definition of the Riemann zeta function $\zeta(s) = \sum_{n>0} \frac{1}{n^s}$ for $s \in \mathbb{C}$.
- Unlike the Riemann zeta function, however, we can actually just evaluate the zeta function for A: since there are q^d monic polynomials of degree d, we see that $\sum_{\deg(f)\leq d \text{ monic}} \frac{1}{|f|^s} = 1 + \frac{q}{q^s} + \frac{q^2}{q^{2s}} + \cdots + \frac{q^d}{q^{ds}} = \frac{1-q^{(d+1)(1-s)}}{1-q^{1-s}}$, and so taking $d \to \infty$ we see that $\zeta_A(s) = \frac{1}{1-q^{1-s}}$ whenever $\mathrm{Re}(s) > 1$ (to ensure convergence).
- We have an obvious meromorphic continuation for $\zeta_A(s)$ to the complex plane (i.e., via the formula above), and it is clear that ζ is analytic everywhere except for a simple pole at $s = 1$.
- Exercise: Show that the residue of $\zeta_A(s)$ at $s = 1$ (which is to say, the value of $\lim_{s\to 1}(s-1)\zeta_A(s)$) is $1/\log q$.
- We also have a functional equation for $\zeta_A(s)$: if we set $\xi_A(s) = q^{-s}(1-q^{-s})^{-1}\zeta_A(s)$, then $\xi_A(s) = \xi_A(1-s)$.
- Exercise: Do the algebra to establish the functional equation.

- We can also represent $\zeta_A(s)$ as an Euler product, just as with the Riemann zeta function.
  - Explicitly, by the uniqueness of prime factorization, we can formally write $\zeta_A(s) = \sum_{f\in A \text{ monic}} \frac{1}{|f|^s} = \prod_{p\text{monic irred}}(1 + \frac{1}{|p|^s} + \frac{1}{|p|^{2s}} + \cdots) = \prod_{p\text{monic irred}}(1 - 1/|p|^s)^{-1}$, and both sides are absolutely convergent for $\mathrm{Re}(s) > 1$.
  - To prove this equality rigorously, we need to do some estimations on tails of the respective series, but since everything converges absolutely, this is not so difficult; we leave the precise details as an exercise.
- We can use the Euler product for the zeta function to obtain the same prime counts that we got earlier.
- Proposition (Prime Counting, Again): If a_d is the number of irreducible monic polynomials in $A = \mathbb{F}_q[t]$ of degree d, then $\sum_{d|n} da_d = q^n$, and so by Mobius inversion as before, we see $a_n = \frac{1}{n}\sum_{d|n} \mu(d) q^{n/d}$.
  - Proof: Group the terms in the Euler product together by degree: if $\deg p = d$ then $|p|^s = q^{ds}$.
  - Thus, since there are a_d monic irreducibles of degree d by definition, we see that $\zeta_A(s) = \prod_{p\text{monic irred}}(1 - 1/|p|^s)^{-1} = \prod_{d=1}^{\infty}(1-q^{-ds})^{-a_d}$.
  - Noting from earlier that $\zeta_A(s) = \frac{1}{1-q^{1-s}}$, if we substitute $u = q^{-s}$, we obtain the equality $\frac{1}{1-qu} = \prod_{d=1}^{\infty}(1-u^d)^{-a_d}$.
  - Taking the log-derivative of both sides yields $\frac{q}{1-qu} = \sum_{d=1}^{\infty} \frac{da_d u^{d-1}}{1-u^d}$. These expressions are equal as power series in u, and thus corresponding coefficients must also be equal.
  - The LHS is $\frac{q}{1-qu} = q\sum_{k=0}^{\infty}(qu)^k$ while the RHS is $\sum_{d=1}^{\infty} da_d u^{d-1} \sum_{l=0}^{\infty} u^{dl} = \sum_{d=1}^{\infty}\sum_{l=0}^{\infty} da_d u^{d(l+1)-1}$. So the coefficient of u^{n-1} on the LHS is $q \cdot q^{n-1} = q^n$, while the coefficient of u^{n-1} on the RHS is $\sum_{d(l+1)=n} da_d = \sum_{d|n} da_d$.
  - Thus, $q^n = \sum_{d|n} da_d$ as claimed.

- Of course, we have already proven this result by counting irreducible polynomials algebraically. However, this approach using the zeta function also extends to solve other counting problems quite conveniently.
- Proposition (Squarefree Counting): The number of monic squarefree polynomials of degree n over $\mathbb{F}_q[t]$ is equal to $b_n := q^n - q^{n-1}$. Equivalently, a randomly-chosen degree-n polynomial is squarefree with probability $1 - 1/q = 1/\zeta_A(2)$.
  - Compare this result to the corresponding fact about integers (which is a little harder to pose because we have to phrase it over a range): if α_n is the probability that a randomly-chosen integer in $[1, n]$ is squarefree, then $\lim_{n\to\infty} \alpha_n = 6/\pi^2 = 1/\zeta(2)$.

- Proof: Consider the product $\pi = \prod_{p \text{ monic irred}}(1 + |p|^{-s})$.
- By multiplying out the terms, we see that for $\text{Re}(s) > 1$, we have $\pi = \sum_{f \text{ monic}} \frac{\delta(f)}{|f|^s}$ where $\delta(f) = \begin{cases} 1 & \text{if } f \text{is squarefree} \\ 0 & \text{if } f \text{is not squarefree} \end{cases}$, since the denominators in the Euler product only include prime factors of exponents 0 and 1.
- Now, since $1 + |p|^{-s} = \frac{1 - |p|^{-2s}}{1 - |p|^{-s}}$, taking the product over monic irreducibles and using the fact that the resulting numerator and denominator products converge absolutely allows us to write $\pi = \prod_{p \text{ monic irred}} \frac{1 - |p|^{-2s}}{1 - |p|^{-s}} = \frac{\prod_{p \text{ monic irred}} 1 - |p|^{-2s}}{\prod_{p \text{ monic irred}} 1 - |p|^{-s}} = \frac{\zeta_A(2s)}{\zeta_A(s)}$.
- Setting $u = q^{-s}$ yields $\frac{1 - qu^2}{1 - qu} = \frac{\zeta_A(2s)}{\zeta_A(s)} = \pi = \sum_{f \text{ monic}} \frac{\delta(f)}{|f|^s} = \sum_{n=0}^{\infty} b_n u^n$.
- But as a power series in u, we have $\frac{1 - qu^2}{1 - qu} = (1 - qu^2)(1 + qu + q^2u^2 + \cdots)$, and so comparing coefficients yields $b_n = q^n - q^{n-1}$ as claimed.

- In a similar way, we can use the zeta function to write down formulas for the number of monic kth-powerfree polynomials of a given degree over $\mathbb{F}_q[t]$.
- Specifically, these values are packaged as the coefficients in the Euler product $\prod_{p \text{ monic irred}}(1 + |p|^{-s} + |p|^{-2s} + \cdots + |p|^{-(k-1)s}) = \frac{\zeta_A(ks)}{\zeta_A(s)}$, and then by doing a calculation like the one above, one can write down an explicit formula.
- Exercise: Finish this calculation and give the actual formula for the number of cubefree polynomials of degree n.
- It is also worthwhile interpreting this Euler product calculation heuristically in terms of probabilities.
- Explicitly, we would expect (under suitable probability assumptions) that the probability of a given polynomial not being divisible by f is $(1 - 1/|f|)$.
- So, assuming independence (which can be made rigorous by appealing to the Chinese remainder theorem), the probability that a given polynomial is not divisible by any prime power p^k for all monic irreducible p is $\prod_{p \text{ monic irred}}(1 - 1/|p|^k) = 1/\zeta_A(k)$: this is why the $1/\text{zeta}$ factor shows up in the answer.

0.4 (Sep 15) Dirichlet Series and Multiplicative Functions

- Another classical object of study in elementary number theory over $\mathbb{Z}$ are arithmetic functions related to divisors, such as the Euler φ-function, the divisor-counting function, and the sum-of-divisors function.
- All of these are examples of multiplicative functions, which have the property that $f(ab) = f(a)f(b)$ whenever a, b are relatively prime. (Note the infelicitous terminology: if $f(ab) = f(a)f(b)$ for all a, b, f is instead called completely multiplicative.)
- In particular, if n has prime factorization $n = \prod_i p_i^{a_i}$ and f is multiplicative, then $f(n) = \prod_i f(p_i^{a_i})$.
- We will briefly review some results about multiplicative functions in the classical setting, and then redo them in the function-field setting.

- It is a standard combinatorial principle that if we want to understand a function with domain $\mathbb{N}$, we should look at its generating function.
- A natural first guess would be to use the standard power series $\sum_{n=0}^{\infty} f(n)x^n$.
- However, this type of generating function is useful primarily for functions that behave additively. For number-theoretic functions, we instead want to use a Dirichlet series.

- Definition: If $h : \mathbb{N} \to \mathbb{C}$ is a complex-valued function defined on positive integers, then its associated Dirichlet series is $D_h(s) = \sum_{n=1}^{\infty} \frac{h(n)}{n^s}$.
  - Example: If $h(n) = 1$ for all n, then $D_h(s) = \zeta(s)$, the Riemann zeta function.
  - In order for this series to converge, we need h not to grow too fast. One may check that if $h(n) = O(n^\alpha)$ then $D_h(s)$ is absolutely convergent for $\mathrm{Re}(s) > 1 + \alpha$. (We will mostly ignore issues of convergence, since our functions will grow polynomially at worst, and so we may manipulate the series as if they were formal power series.)
  - If h is multiplicative, then it is a straightforward calculation to see that $D_h(s)$ has an Euler product expansion: $D_h(s) = \prod_{p \text{ prime}} (1 + \frac{h(p)}{p} + \frac{h(p^2)}{p^2} + \cdots)$, on the appropriate domain of convergence.
- The key property of Dirichlet series is that they reproduce desired behaviors under multiplication:
- Proposition (Dirichlet Multiplication): If $f, g : \mathbb{N} \to \mathbb{C}$ are functions, then $D_f(s) \cdot D_g(s) = D_{f*g}(s)$ where $f * g$ is the Dirichlet convolution defined via $(f * g)(n) = \sum_{d|n} f(d)g(n/d)$.
  - Proof: $D_f(s)D_g(s) = \sum_{a=1}^{\infty} \sum_{b=1}^{\infty} \frac{f(a)g(b)}{(ab)^s} = \sum_{n=1}^{\infty} \frac{1}{n^s} \sum_{ab=n} f(a)g(b) = \sum_{n=1}^{\infty} \frac{(f*g)(n)}{n^s} = D_{f*g}(s)$.
- The Dirichlet convolution, owing to the fact that it is merely multiplication of the underlying Dirichlet series, has various nice properties.
  - Exercise: Show that Dirichlet convolution is commutative and associative, and has an identity element given by $I(n) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$.
  - Exercise: Show that f has an inverse under Dirichlet convolution if and only if $f(1) \neq 0$.
  - Exercise: If $f(1) \neq 0$ and f is multiplicative, then its Dirichlet inverse f^{-1} is also multiplicative.
  - Exercise: Show that if two of f, g, and $f * g$ are multiplicative, then the third is also.
- By exploiting Dirichlet convolution, we can find the Dirichlet series for many basic multiplicative functions in terms of the Riemann zeta function.
  - Recall $I(n) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$ and the Mobius function $\mu(n) = \begin{cases} 0 & \text{if } n \text{is not squarefree} \\ (-1)^r & \text{if } n \text{is the product of } r \text{distinct primes} \end{cases}$.
  - Also define $N(n) = n$ and $1(n) = 1$ (for all n).
  - Exercise: Show that $D_I(s) = 1$, $D_1(s) = \zeta(s)$, and $D_N(s) = \zeta(s-1)$.
  - First, we note that $\mu * 1 = I$, since $(\mu * 1)(n) = \sum_{d|n} \mu(d)1(n/d) = \sum_{d|n} \mu(d) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$ as noted in an exercise previously. Therefore, by multiplicativity of the Dirichlet series, we see that $D_\mu(s)D_1(s) = D_I(s)$, so that $D_\mu(s) = \frac{1}{\zeta(s)}$.
  - Exercise: Use $\mu * 1 = I$ to establish Mobius inversion: if $g(n) = \sum_{d|n} f(n)$ then $f(n) = \sum_{d|n} \mu(d)g(n/d)$.
  - Exercise: For the Euler φ-function, show that $\sum_{d|n} \varphi(d) = n$.
  - The previous exercise says that $\varphi * 1 = N$, and so by composing with μ and using associativity, we see that $\varphi = \mu * N$. Then we have $D_\varphi(s) = D_\mu(s)D_N(s) = \frac{\zeta(s-1)}{\zeta(s)}$.
  - In principle, we could have established this fornula for $D_\varphi(s)$ by manipulating the zeta function directly, but this method is both more difficult and requires knowing the actual (non-obvious) formula for the answer ahead of time.
  - We can also find the Dirichlet series for the divisor-counting function $d(n) = \#\{d \in \mathbb{N} : d|n\}$ quite easily by noting that $d(n) = \sum_{d|n} 1(d)1(d/n)$: this means $d = 1 * 1$, so $D_d(s) = D_1(s)^2 = \zeta(s)^2$.

- Exercise: If σ is the sum-of-divisors function $\sigma(n) = \sum_{d|n} d$, show that $D_\sigma(s) = \zeta(s)\zeta(s-1)$.
- Exercise: If σ_k is the sum-of-kth-powers-of-divisors function $\sigma_k(n) = \sum_{d|n} d^k$, find and prove a formula for $D_{\sigma_k}(s)$ in terms of the Riemann zeta function.

- One of the main applications of computing the Dirichlet series for these various arithmetic functions is that we can extract information about average growth rates from them.
  - In the classical case, obtaining average-growth results is moderately delicate, so we will instead just focus on the function-field case.

- Here are the function-field analogues of these classical multiplicative functions, which are now complex-valued functions on monic polynomials rather than positive integers:
  - The identity: $I(f) = \begin{cases} 1 & \text{for } f = 1 \\ 0 & \text{for } f \neq 1 \end{cases}$.
  - The norm: $N(f) = |f|$.
  - The Mobius μ-function: $\mu(f) = \begin{cases} 0 & \text{if } f \text{is not squarefree} \\ (-1)^r & \text{if } f \text{is the product of } r \text{distinct primes} \end{cases}$.
  - The Euler Φ-function: $\Phi(f) = \#(A/fA)^* = |f| \prod_{p|f \text{ prime}} (1 - 1/|p|)$.
  - The divisor-counting function: $d(f) = \#\{\text{monic } d|f\}$.
  - The sum-of-divisors function: $\sigma(f) = \sum_{d|f \text{ monic}} |d|$, or more generally the sum-of-kth-powers-of-divisors function $\sigma_k(f) = \sum_{d|f \text{ monic}} |d|^k$. (Note here that we take the norm of the divisors, since we want a $\mathbb{C}$-valued function.)
  - It is easy to check that all of these functions are multiplicative, and to write down formulas for all of them in terms of the prime factorization of $f = p_1^{a_1} \cdots p_k^{a_k}$.
  - Exercise: Verify that $d(f) = (a_1 + 1) \cdots (a_k + 1)$ and $\sigma(f) = \dfrac{|p_1|^{a_1+1} - 1}{|p_1| - 1} \cdots \dfrac{|p_k|^{a_k+1} - 1}{|p_k| - 1}$.

- We have essentially the same definition for the Dirichlet series in the function-field case:
- Definition: If $h : \{\text{monics}\} \to \mathbb{C}$ is a complex-valued function defined on monic polynomials in $\mathbb{F}_q[t]$, then its associated Dirichlet series is $D_h(s) = \sum_{f \text{ monic}} \dfrac{h(f)}{|f|^s}$.
  - As before, we will mostly ignore issues of convergence, but just as in the classical case, one may check that if $h(f) = O(|f|^\alpha)$ then $D_h(s)$ converges absolutely for $\text{Re}(s) > 1 + \alpha$.
  - We also have the same Dirichlet convolution operator: if $g, h : \{\text{monics}\} \to \mathbb{C}$ are functions, then $D_g(s) \cdot D_h(s) = D_{g \star h}(s)$ where $(g * h)(f) = \sum_{d|f \text{ monic}} g(d)h(f/d)$.
  - Dirichlet convolution is commutative, associative, and has the identity element $I(f) = \begin{cases} 1 & \text{for } f = 1 \\ 0 & \text{for } f \neq 1 \end{cases}$.
  - All of the same formulas for our arithmetic functions in terms of the zeta function follow through just as before. Here, however, we can actually write out the expressions explicitly, since we have a formula $\zeta_A(s) = \dfrac{1}{1 - q^{1-s}}$.

- Proposition (Some Dirichlet Series): For $u = q^{-s}$, we have the following formulas: $D_I(s) = 1$, $D_N(s) = \zeta_A(s-1) = \dfrac{1}{1-u}$, $D_1(s) = \zeta_A(s) = \dfrac{1}{1-qu}$, $D_\mu(s) = \dfrac{1}{\zeta_A(s)} = 1 - qu$, $D_\Phi(s) = \dfrac{\zeta_A(s-1)}{\zeta_A(s)} = \dfrac{1-qu}{1-q^2u}$, $D_d(s) = \zeta_A(s)^2 = \dfrac{1}{(1-qu)^2}$, and $D_\sigma(s) = \zeta_A(s)\zeta_A(s-1) = \dfrac{1}{(1-qu)(1-q^2u)}$.
  - Proof: Exercise.

- Using these formulas we can recover average-value results quite easily.
- Definition: If $h : \{\text{monics}\} \to \mathbb{C}$ is a function, the average value of h on degree-n polynomials is $\text{Avg}_n(h) = \frac{1}{q^n} \sum_{\deg(f)=n \text{ monic}} h(f)$. If the limit $\lim_{n\to\infty} \text{Avg}_n(h)$ exists, we call it the "average value" of h.
  - We can also easily average h on polynomials of degree $\leq n$: the desired sum is instead $\frac{1}{1+q+\cdots+q^n} \sum_{\deg(f)\leq n} h(f)$.
  - Exercise: Show that if $\lim_{n\to\infty} \text{Avg}_n(h) = \alpha$, then $\lim_{n\to\infty} \frac{1}{1+q+\cdots+q^n} \sum_{\deg(f)\leq n} h(f) = \alpha$ as well, so it is irrelevant whether we average over degree exactly n or $\leq n$.
  - The nice result here is that we can read off the value of $\text{Avg}_n(h)$ from the coefficients of the Dirichlet series for h: explicitly, we have $D_h(s) = \sum_{n=1}^{\infty} \frac{\sum_{\deg(f)=n} h(f)}{q^{ns}} = \sum_{n=1}^{\infty} \frac{q^n \text{Avg}_n(h)}{q^{ns}} = \sum_{n=1}^{\infty} q^n \text{Avg}_n(h) u^n$ for $u = q^{-s}$.
  - So we can calculate these averages by simply expanding out the Dirichlet series calculated above as power series in $u = q^{-s}$ and then dividing by q^n.
  - For example, $D_\mu(s) = 1 - qu$, so the average value of μ is 1 on degree-0 polynomials, -1 on degree-1 polynomials, and 0 on higher-degree polynomials.
  - Similarly, $D_d(s) = \frac{1}{(1-qu)^2} = (1 + qu + q^2u^2 + \cdots)^2 = 1 + 2qu^2 + 3q^2u^3 + \cdots$, so the average value of d on degree-n polynomials is $n+1$.
  - Likewise, $D_\Phi(s) = \frac{1-qu}{1-q^2u} = (1-qu)(1 + q^2u + q^4u^2 + q^6u^3 + \cdots) = 1 + (q^2-q)u + (q^4-q^3)u^2 + \cdots$, so the average value of Φ on degree-n polynomials is $(q^{2n} - q^{2n-1})/q^n = q^n - q^{n-1}$.
  - Exercise: Show that the average value of σ on degree-n polynomials is $(q^{n+1} - 1)/(q-1)$.

0.5 (Sep 17) Primes in Arithmetic Progressions, Part 1

- Our next task is to prove the function-field analogue of Dirichlet's theorem on primes in arithmetic progressions.
  - Over $\mathbb{Q}$, Dirichlet's theorem says that for any positive integer m and any a relatively prime to m, there exist infinitely many primes in the arithmetic progression $\{a, a+m, a+2m, a+3m, \dots\}$: in other words, congruent to a modulo m.
  - Exercise (easy): Show that if a is not relatively prime to m, then there are only finitely many primes congruent to a modulo m.
- There are $\varphi(m)$ residue classes modulo m that contain infinitely many primes, so one can ask more precisely about how the primes are distributed among these residue classes.
  - In fact, the primes are asymptotically uniformly distributed among these residue classes: the proportion of primes congruent to a modulo m approaches $1/\varphi(m)$ upon taking an appropriate limit.
  - Explicitly, define the natural density of a set S of primes to be $\lim_{n\to\infty} \frac{S \cap \{1,2,\dots,n\}}{\{\text{primes}\} \cap \{1,2,\dots,n\}}$, provided the limit exists.
  - Then, as first proven by de la Vallée Poussin, the natural density of the primes congruent to a modulo m is $1/\varphi(m)$ when a is relatively prime to m.
- However, the natural density is somewhat difficult to handle with analytic methods. From the standpoint of zeta functions, a more natural choice is the Dirichlet density:
- Definition: If S is a set of primes, the Dirichlet density of S is the value $\delta_S = \lim_{s\to 1+} \frac{\sum_{\text{primes } p\in S} p^{-s}}{\sum_{\text{primes } p} p^{-s}}$, assuming the limit exists.

- Note that the sum in the numerator is always finite for $\mathrm{Re}(s) > 1$ by comparison to the sum for the zeta function.
- Exercise: If S is finite, show that its Dirichlet density is 0.
- One may prove that if a set has natural density δ, then its Dirichlet density is also δ. The converse is not true, however: a simple counterexample due to Serre is the set S of primes whose leading digit is 1 in base 10.
- Exercise (hard): Show that the set of primes whose leading digit is 1 in base 10 has undefined natural density, but has Dirichlet density $\log_{10} 2$. (The answer works out the same if you use integers with leading digit 1.)

- The corresponding definition for function fields is as follows:

- Definition: If T is a set of monic irreducibles in $\mathbb{F}_q[t]$, its Dirichlet density is $\delta_T = \lim_{s \to 1+} \frac{\sum_{p \in T} |p|^{-s}}{\sum_p |p|^{-s}}$, assuming the limit exists.
  - We note that both the numerator and denominator sums converge for $\mathrm{Re}(s) > 1$.

- Our main result is the following:

- Theorem (Analogue of Dirichlet's Theorem): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the Dirichlet density of the set of primes congruent to $a \pmod m$ exists and is $1/\Phi(m)$. In particular, there are infinitely many such primes.
  - The fundamentally hard part of proving this theorem is to establish the nonvanishing of the L-functions for nontrivial characters at $s = 1$.
  - In order to explain what this means (and then do it), we will begin with a brisk discussion of Dirichlet characters and their properties.

- Definition: Let G be a finite abelian group. A group character χ of G is a homomorphism $\chi : G \to \mathbb{C}^\times$.
  - Note that $\chi(1) = 1$ for every character, and also if $g \in G$ has order d, then $1 = \chi(1) = \chi(g^d) = \chi(g)^d$, so $\chi(g)$ is a dth root of unity. Thus in general, χ is a map from G to the group of complex $|G|$th roots of unity.
  - Example: For any G, the trivial character χ_{triv} has $\chi_{\mathrm{triv}}(g) = 1$ for all $g \in G$.
  - Example: If $G = (\mathbb{Z}/p\mathbb{Z})^\times$, the quadratic residue symbol $\chi(a) = \left(\frac{a}{p}\right)$ is a group character.
  - Example: If $G = (A/pA)^\times$ for $A = \mathbb{F}_q[t]$ and d divides $q-1$, the dth-power residue symbol $\chi(a) = \left(\frac{a}{p}\right)_d$ gives a group character, provided we identify the dth roots of unity in $\mathbb{F}_q$ with the dth roots of unity in $\mathbb{C}$ (simply choose any fixed isomorphism).

- We will be interested in the case where G is the group of units $(\mathbb{Z}/m\mathbb{Z})^\times$ or $(A/fA)^\times$, in which case we call χ a Dirichlet character.
  - In some situations it is slightly more convenient to work with extended Dirichlet characters, which we extend to have domain $\mathbb{Z}/m\mathbb{Z}$ or A/fA by setting $\chi(a) = 0$ whenever a is not relatively prime to the modulus.
  - Exercise: Extended Dirichlet characters modulo m are the same as functions $\chi : \mathbb{Z} \to \mathbb{C}$ (or $A \to \mathbb{C}$) such that (i) $\chi(a + bm) = \chi(a)$ for all a, b, (ii) $\chi(ab) = \chi(a)\chi(b)$ for all a, b, and (iii) $\chi(a) \neq 0$ iff a is relatively prime to m.

- We can multiply two group characters on G pointwise, and this operation makes them into a group:

- Proposition (Dual Group of G): The set of group characters on G forms a group under pointwise multiplication. The identity is the trivial character and the inverse of χ is its complex conjugate $\overline{\chi}$. This group is called the dual group of G and is denoted $\hat{G}$.

- Proof: These properties can be checked directly (exercise), or one may simply note that $\hat{G} = \mathrm{Hom}(G, \mathbb{C}^\times)$.

- The dual group $\hat{G}$ is also an abelian group, so it is natural to wonder how its structure relates to G. In fact, it is isomorphic to G:

- Proposition (Dual Group, II): If G is a finite abelian group, its dual group $\hat{G}$ is isomorphic to G.

  - Proof: First consider the special case where G is a cyclic group of order n generated by g. Then $\chi(g^d) = \chi(g)^d$ for all d, so any group character χ is uniquely determined by the value of $\chi(g)$, which must be some nth root of unity.
  - Conversely, any such selection $e^{2\pi ia/n}$ for $\chi(g)$ yields a valid group character χ_a, namely with $\chi_a(g^d) = e^{2\pi iad/n}$. Since $\chi_a\chi_b = \chi_{a+b}$ and χ_1^n is the trivial character, we see that the dual group $\hat{G}$ is cyclic of order n (the map $a \mapsto \chi_a$ is an isomorphism of $\hat{G}$ with $\mathbb{Z}/n\mathbb{Z}$).
  - Now suppose $G = H \times K$ is a direct product. If $\chi : H \times K \to \mathbb{C}^\times$ is a homomorphism, let $\chi_H : H \to \mathbb{C}^\times$ and $\chi_K : K \to \mathbb{C}^\times$ be the projections $\chi_H(h) = \chi(h, 1)$ and $\chi_K(k) = \chi(1, k)$. Then χ_H is a group character of H, χ_K is a group character of K, and $\chi = \chi_H\chi_K$.
  - Conversely, any pair $(\chi_H, \chi_K) \in (\hat{H}, \hat{K})$ yields a character $\chi = \chi_H\chi_K \in \hat{G}$, so we see $\hat{G} \cong \hat{H} \times \hat{K}$.
  - Since every finite abelian group is a direct product of cyclic groups, and the result holds for cyclic groups and direct products, we are done.

- Exercise: If H is a subgroup of the finite abelian group G, define $H^\perp = \{\chi \in \hat{G} : \chi(H) = 1\}$. Show that $H^\perp \cong \widehat{G/H}$ and that $\hat{G}/H^\perp \cong \hat{H}$. Use these results along with $\hat{G} \cong G$ to conclude that the subgroup lattice of G is the same when turned upside down.

- The isomorphism between $\hat{G}$ and G above is non-canonical (i.e., it is not "coordinate-free" in the sense that we must pick specific generators for G and $\hat{G}$ to obtain the isomorphism).

  - However, there is a canonical isomorphism between $\hat{\hat{G}}$ (the double dual) and G given by the "evaluation map" φ, which maps an element $g \in G$ to the "evaluation-at-g" map e_g on characters $\chi \in \hat{G}$, defined by $e_g(\chi) = \chi(g)$.
  - Exercise: Verify that the evaluation map $\varphi : G \to \hat{\hat{G}}$ with $\varphi(g) = \{\chi \mapsto \chi(g)\}$ is an isomorphism from $\hat{\hat{G}}$ to G.
  - This result is a special case of Pontryagin duality, and has an analogous statement for duals of finite-dimensional vector spaces.
  - In fact, it is really the algebraic analogue of Fourier inversion (the reason being that Fourier analysis on finite abelian groups involves sums over group characters in lieu of integrals). For a brief taste of the analogy, the main idea is to note that the map $e^{inx} : \mathbb{R} \to \mathbb{C}^\times$ is a group homomorphism, and thus is an "$\mathbb{R}$"-character.

- We can also put the structure of an inner product on group characters. To establish this we first show some simple orthogonality relations:

- Proposition (Orthogonality Relations): If G is a finite abelian group and χ is a group character, the following hold:

  1. The sum $\sum_{g \in G} \chi(g) = \begin{cases} |G| & \text{if } \chi\text{is trivial} \\ 0 & \text{otherwise} \end{cases}$.
     - Proof: If χ is trivial the sum is clearly $|G|$. If χ is not trivial, say with $\chi(h) \neq 1$, then $\sum_{g \in G} \chi(g) = \sum_{g \in G} \chi(gh) = \chi(h) \sum_{g \in G} \chi(g)$ by reindexing (since $G = Gh$), and so $\sum_{g \in G} \chi(g) = 0$.
  2. The sum $\sum_{\chi \in \hat{G}} \chi(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{otherwise} \end{cases}$.
     - Proof: Apply Pontryagin duality to (1).

3. (Orthogonality 1) For any characters χ_1 and χ_2, $\sum_{g\in G}\chi_1(g)\overline{\chi_2(g)} = \begin{cases} |G| & \text{if } \chi_1 = \chi_2 \\ 0 & \text{otherwise} \end{cases}$.
   - Proof: Apply (1) to $\chi = \chi_1\overline{\chi_2}$.
4. (Orthogonality 2) For any elements g_1 and g_2, $\sum_{\chi\in\hat{G}}\chi(g_1)\overline{\chi(g_2)} = \begin{cases} |G| & \text{if } g_1 = g_2 \\ 0 & \text{otherwise} \end{cases}$.
   - Proof: Apply (2) to $g = g_1 g_2^{-1}$, or apply Pontryagin duality to (3).
5. The pairing $\langle f_1, f_2\rangle_G = \frac{1}{|G|}\sum_{g\in G} f_1(g)\overline{f_2(g)}$ is a complex inner product on functions $f : G \to \mathbb{C}$, and the elements of the dual group $\hat{G}$ are an orthonormal basis with respect to this inner product.
   - Proof: The inner product axioms are straightforward, and the fact that $\hat{G}$ yields an orthonormal basis follows from (3).
6. The pairing $\left\langle \hat{f}_1, \hat{f}_2\right\rangle_{\hat{G}} = \frac{1}{|G|}\sum_{\chi\in\hat{G}} \hat{f}_1(\chi)\overline{\hat{f}_2(\chi)}$ is a complex inner product on functions $\hat{f} : \hat{G} \to \mathbb{C}$, and the elements of G are an orthonormal basis with respect to this inner product.
   - Proof: The inner product axioms are straightforward, and the fact that $G \cong \hat{\hat{G}}$ yields an orthonormal basis follows from (4), or apply Pontryagin duality to (5).
7. (Fourier Inversion) For any function $f : G \to \mathbb{C}$, with the Fourier transform $\hat{f} : \hat{G} \to \mathbb{C}$ defined by $\hat{f}(\chi) = \langle f, \chi\rangle_G = \frac{1}{|G|}\sum_{g\in G} f(g)\overline{\chi(g)}$, we have $f(g) = \sum_{\chi\in\hat{G}}\hat{f}(\chi)\chi(g)$ for all $g \in G$.
   - Proof: This follows immediately from (5), since the elements of $\hat{G}$ are an orthonormal basis.

- Exercise: Prove Plancherel's theorem $\langle f_1, f_2\rangle_G = \frac{1}{|G|}\left\langle \hat{f}_1, \hat{f}_2\right\rangle_{\hat{G}}$ and deduce Parseval's theorem $\sum_{g\in G}|f(g)|^2 = \frac{1}{|G|}\sum_{\chi\in\hat{G}}\left|\hat{f}(\chi)\right|^2$.
- With the fundamentals taken care of, we can now focus on Dirichlet characters.
   - Studying primes congruent to a modulo m naturally leads to a question about Dirichlet characters via Fourier inversion, since we may decompose the characteristic function of [primes congruent to a modulo m] as a sum over Dirichlet characters for the group $G = (A/mA)^*$.
   - Explicitly, if $\delta_a(p)$ is 1 when $p \equiv a \pmod m$ and 0 otherwise, then $\hat{\delta_a}(\chi) = \frac{1}{\Phi(m)}\sum_{g\in G}\delta_a(g)\overline{\chi(g)} = \frac{1}{\Phi(m)}\overline{\chi(a)}$, since the only nonzero value of $\delta_a(g)$ occurs when $g \equiv a \pmod m$.
   - Then by Fourier inversion we have $\delta_a(p) = \sum_{\chi\in\hat{G}}\hat{\delta_a}(\chi)\chi(p) = \sum_{\chi\in\hat{G}}\frac{1}{\Phi(m)}\overline{\chi(a)}\chi(p)$. So the numerator for the Dirichlet density is $\sum_{p\equiv a \pmod m}|p|^{-s} = \sum_p \delta_a(p)\,|p|^{-s} = \frac{1}{\Phi(m)}\sum_{\chi\in\hat{G}}\left[\overline{\chi(a)}\sum_p \chi(p)\,|p|^{-s}\right]$.
   - This is a bit complicated, but the point is that we have a sum over the Dirichlet characters of constants (namely $\overline{\chi(a)}$) times $\sum_p \frac{\chi(p)}{|p|^s}$, which is quite close to the Dirichlet series for the character χ (the only difference is that we are only summing over primes, rather than all monic polynomials).
   - As we will see, we will be able to extract this sum over primes from the full Dirichlet series, which we now examine more closely.
   - The main reason we go to this effort to use Fourier inversion is that the Dirichlet series for Dirichlet characters behave very nicely (far more nicely than the original series over primes congruent to a modulo m) because Dirichlet characters are multiplicative.

- Definition: If χ is a Dirichlet character modulo m, we define its associated Dirichlet L-series $L(s,\chi) = \sum_{f \text{ monic}}\frac{\chi(f)}{|f|^s}$.

- Note that this is just the Dirichlet series for $\chi(f)$, as we defined it previously. It is traditional to denote these series with the letter L (which was the letter Dirichlet used for such functions).
- As usual, the series converges absolutely for $\mathrm{Re}(s) > 1$, since $|\chi(f)| \leq 1$ for all f.
- Furthermore, because Dirichlet characters are completely multiplicative, the L-series has a very simple Euler product: explicitly, $L(s,\chi) = \prod_{p \text{ irred}} \left[1 - \frac{\chi(p)}{|p|^s}\right]^{-1}$, for $\mathrm{Re}(s) > 1$.
- The Euler product is the key to calculating the Dirichlet density we wanted earlier: taking the logarithm of the Euler product gives $\log L(s,\chi) = -\sum_{p \text{ irred}} \log(1 - \chi(p)/|p|^s) \approx \sum_{p \text{ irred}} \frac{\chi(p)}{|p|^s}$ using the Taylor approximation $-\log(1-x) \approx x$ which is accurate for small $|x|$.
- So our main task is to determine what happens to $\log L(s,\chi)$ as $s \to 1$, since this is the required input for calculating the Dirichlet density of the primes congruent to a modulo m.

0.6 (Sep 22) Primes in Arithmetic Progressions, Part 2

- Our main task is to determine what happens to $\log L(s,\chi)$ as $s \to 1$, since this is the required input for calculating the Dirichlet density of the primes congruent to a modulo m.
- Example: For the trivial character χ_{triv}, we have $L(s,\chi_{\text{triv}}) = \prod_{p|m \text{ irred}}(1 - |p|^{-s}) \cdot \zeta_A(s)$, since the terms with $p|m$ are missing from the Euler product for $L(s,\chi)$.
  - In particular, we see that $L(s,\chi_{\text{triv}})$ has an analytic continuation (since $\zeta_A(s)$ does) and a single simple pole at $s = 1$.
- For other characters, the L-series is essentially finite.
- Proposition (L-Series for Nontrivial Characters): Let m be a monic polynomial of positive degree and χ be a nontrivial Dirichlet character modulo m. Then $L(s,\chi)$ is a polynomial in q^{-s} of degree at most $\deg m - 1$, and in particular has an analytic continuation.
  - Proof: Let $A(n,\chi) = \sum_{\deg f = n} \chi(f)$ and note, as we have previously done in working out average-value results, that $L(s,\chi) = \sum_{n=0}^{\infty} A(n,\chi) q^{-ns}$. The claimed result is then equivalent to saying $A(n,\chi) = 0$ for $n \geq \deg m$.
  - For this, suppose $\deg f = n \geq m$ and write $f = hm + r$ with $\deg r < \deg m$, where $\deg h = \deg f - \deg m$ and $\mathrm{sgn}(h) = 1/\mathrm{sgn}(m)$. Conversely, given such an h and r, we get a unique $f = hm + r$. Note that $\chi(f) = \chi(r)$, and also that there are $q^{n - \deg m}$ possible h.
  - Then $A(n,\chi) = \sum_{\deg f = n} \chi(f) = \sum_{\deg f = n} \chi(r) = q^{n-\deg m} \sum_{\deg r < \deg m} \chi(r) = 0$ where the last sum is zero by the orthogonality relation (1).
  - The observation about the analytic continuation is immediate (simply take the analytic continuation as the given polynomial in q^{-s}).
- Exercise: Choose a modulus $m \in \mathbb{F}_q[t]$ and a nontrivial Dirichlet character χ, and verify explicitly that $L(s,\chi)$ is a polynomial in q^{-s}.
- As a consequence, we see that $L(s,\chi)$ has no pole at $s = 1$ when $\chi \neq \chi_{\text{triv}}$. Our next major goal is to prove that $L(1,\chi) \neq 0$ for $\chi \neq \chi_{\text{triv}}$.
- Lemma: Let χ be any Dirichlet character modulo m. Then for each monic irreducible p not dividing m, there exist $f_p, g_p > 0$ with $f_p g_p = \Phi(m)$ such that $\prod_{\chi \in \hat{G}} L(s,\chi) = \prod_{p \nmid m}(1 - |p|^{-f_p s})^{-g_p}$.
  - Proof: For a fixed monic irreducible $p \nmid m$, as we have previously noted the evaluation-at-p map $\chi \mapsto \chi(p)$ is a homomorphism from $\hat{G}$ to $\mathbb{C}^\times$.
  - Let the image be a cyclic group of order f_p and the kernel have size g_p: then $f_p g_p = \#\hat{G} = \#G = \Phi(m)$ by the first isomorphism theorem.

- For this p, by grouping the fibers of the evaluation-at-p map together, for $\zeta = e^{2\pi i/f_p}$ we have $\prod_{\chi\in\hat{G}}(1-\chi(p)/|p|^s)^{-1} = \prod_{j=0}^{f_p-1}(1-\zeta^j/|p|^s)^{-g_p}$, and this last product equals $(1-|p|^{-f_p s})^{-g_p}$ since it is the evaluation of the polynomial $(1-t)(1-\zeta t)\cdots(1-\zeta^{f_p-1}t) = 1-t^{f_p}$ at $t = |p|^{-s}$.
- Thus, taking the product over all monic irreducibles $p \nmid m$ yields the claimed $\prod_{\chi\in\hat{G}} L(s,\chi) = \prod_{\chi\in\hat{G}}\prod_{p\nmid m}(1-\chi(p)/|p|^s)^{-1} = \prod_{p\nmid m}(1-|p|^{-f_p s})^{-g_p}$ after reversing the order of the products.

- We next show that $L(1,\chi)\neq 0$ for nonreal Dirichlet characters χ:
- Lemma (Nonvanishing, I): Let χ be any Dirichlet character modulo m such that $\chi \neq \overline{\chi}$. Then $L(1,\chi)\neq 0$.
  - Proof: If we expand the product $\prod_{\chi\in\hat{G}} L(s,\chi) = \prod_{p\nmid m}(1-|p|^{-f_p s})^{-g_p}$ from the Lemma above, it yields a Dirichlet series with nonnegative coefficients and constant term 1.
  - Thus, if s is real and greater than 1 (so that the product converges), the value of the product is real and greater than 1.
  - If $\chi \neq \overline{\chi}$, then $\prod_{\chi\in\hat{G}} L(s,\chi) = L(s,\chi_{\text{triv}})L(s,\chi)L(s,\overline{\chi})\cdot[\text{other terms}]$.
  - Now suppose $L(1,\chi) = 0$: then we would have $L(1,\overline{\chi}) = 0$ also. But this would mean the product $\prod_{\chi\in\hat{G}} L(s,\chi)$ vanishes at $s=1$, because the only term that has a pole at $s=1$ is $L(s,\chi_{\text{triv}})$ and that pole has order 1, but we have two zeroes at $s=1$ arising from $L(s,\chi)$ and $L(s,\overline{\chi})$.
  - But this is impossible because the value of the product is real and greater than 1 for $s>1$. Thus, $L(1,\chi)\neq 0$.

- The case where $\chi = \overline{\chi}$ and $\chi \neq \chi_{\text{triv}}$ (i.e., when χ has order 2 in $\hat{G}$) is quite a bit trickier, since we cannot get away with such a simple order-of-vanishing argument.

- Lemma (Nonvanishing, II): Let χ be any Dirichlet character of order 2 modulo m (i.e., such that $\chi = \overline{\chi}$ but $\chi \neq \chi_{\text{triv}}$). Then $L(1,\chi)\neq 0$.
  - Proof: Suppose that $\chi = \overline{\chi}$ but $\chi \neq \chi_{\text{triv}}$, so that $\chi(p)\in\{\pm 1\}$ for $p\nmid m$, and define the function $G(s) = \dfrac{L(s,\chi_{\text{triv}})L(s,\chi)}{L(2s,\chi_{\text{triv}})} = \prod_{p\nmid m}\dfrac{(1-|p|^{-s})^{-1}(1-\chi(p)\,|p|^{-s})^{-1}}{(1-|p|^{-2s})^{-1}} = \prod_{p\nmid m}\dfrac{1+|p|^{-s}}{1-\chi(p)\,|p|^{-s}} = \prod_{p\nmid m,\chi(p)=1}\dfrac{1+|p|^{-s}}{1-|p|^{-s}} = \prod_{p\nmid m,\chi(p)=1}[1+\sum_{k=1}^{\infty}|p|^{-ks}]$.
  - By expanding this last expression for G, we can see that its Dirichlet series has all coefficients nonnegative.
  - We also have $\dfrac{L(s,\chi_{\text{triv}})}{L(2s,\chi_{\text{triv}})} = \dfrac{\zeta_A(s)}{\zeta_A(2s)}\cdot\prod_{p|m}\dfrac{1-|p|^{-s}}{1-|p|^{-2s}} = \dfrac{1-q^{1-2s}}{1-q^{1-s}}\prod_{p|m}(1+|p|^{-s})^{-1}$. Substituting this into the expression for G yields that $\dfrac{1-q^{1-2s}}{1-q^{1-s}}L(s,\chi) = \dfrac{L(s,\chi_{\text{triv}})L(s,\chi)}{L(2s,\chi_{\text{triv}})}\prod_{p|m}(1+|p|^{-s})^{-1} = G(s)\prod_{p|m}(1+|p|^{-s})^{-1}$ is a Dirichlet series with all coefficients nonnegative.
  - Suppose $G(s)\prod_{p|m}(1+|p|^{-s})^{-1} = \sum_{f\text{ monic}}\dfrac{h(f)}{|f|^s}$.
  - Rewriting in terms of $u = q^{-s}$, and noting that $L^*(u,\chi) = L(s,\chi)$ is a polynomial in u as we proved earlier, we obtain the equality $\dfrac{1-qu^2}{1-qu}L^*(u,\chi) = \sum_{d=0}^{\infty}[\sum_{\deg(f)=d}h(f)]u^d$.
  - Now suppose that $L(1,\chi) = L^*(q^{-1},\chi)$ is equal to zero. Then $1-qu$ would divide $L^*(u,\chi)$, which would mean that $\dfrac{1-qu^2}{1-qu}L^*(u,\chi)$ is a polynomial in u. But then the right-hand side would also be a polynomial in u. All of its coefficients are nonnegative (as noted above), which means it cannot have a positive root for u.
  - But, finally, notice that $\dfrac{1-qu^2}{1-qu}L^*(u,\chi)$ is zero when $u = 1/\sqrt{q}$. This is a contradiction, and so $L^*(q^{-1},\chi) = L(1,\chi)$ must be nonzero.

- Now that we know $L(1,\chi)$ vanishes for nontrivial characters χ, we can prove Dirichlet's theorem:

- Theorem (Analogue of Dirichlet's Theorem): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the Dirichlet density of the set of primes congruent to $a \pmod m$ exists and is $1/\Phi(m)$. In particular, there are infinitely many such primes.
  - We have already obtained all of the necessary ingredients, so the proof is mostly a matter of putting them all together.
  - Proof: Recall the power series $-\log(1-x) = \sum_{k=1}^{\infty} x^k/k$, valid for $|x| < 1$.
  - Then for any Dirichlet character χ, we have $\log L(s,\chi) = \sum_p -\log(1-\frac{\chi(p)}{|p|^s}) = \sum_p \left[\sum_{k=1}^{\infty} \frac{\chi(p)^k}{k} |p|^{-ks}\right] = \sum_p \frac{\chi(p)}{|p|^s} + \sum_p \sum_{k=2}^{\infty} \frac{\chi(p)^k}{k} |p|^{-ks}$. The absolute value of the second term is bounded by $\sum_p \sum_{k=2}^{\infty} \frac{1}{k} |p|^{-ks} \leq \sum_{k=2}^{\infty} \sum_{d=1}^{\infty} q^d q^{-kds} \leq \sum_{n=1}^{\infty} (n+1) q^{-ns}$, which is bounded as $s \to 1+$.
  - Therefore, as $s \to 1+$, we have $\log L(s,\chi) = \sum_p \frac{\chi(p)}{|p|^s} + O(1)$. In particular, we see that $\sum_p |p|^{-s} = \log(s-1) + O(1)$ as $s \to 1+$, since $L(s,\chi_{\text{triv}})$ has a simple pole at $s = 1$.
  - Now, by Fourier inversion (as we previously worked out) we have $\sum_{p \equiv a \pmod m} |p|^{-s} = \sum_p \delta_a(p) |p|^{-s} = \frac{1}{\Phi(m)} \sum_{\chi \in \hat{G}} \left[\overline{\chi(a)} \sum_p \frac{\chi(p)}{|p|^s}\right]$.
  - So, the quotient for the Dirichlet density is $\frac{\sum_{p \equiv a \pmod m} |p|^{-s}}{\sum_p |p|^{-s}} = \frac{\frac{1}{\Phi(m)} \sum_{\chi \in \hat{G}} \left[\overline{\chi(a)} \sum_p \frac{\chi(p)}{|p|^s}\right]}{\sum_p |p|^{-s}} = \frac{1}{\Phi(m)} \left[\frac{\sum_{p \nmid m} |p|^{-s}}{\sum_p |p|^{-s}} + \frac{\sum_{\chi \neq \chi_{\text{triv}}} \overline{\chi(a)} \sum_p \frac{\chi(p)}{|p|^s}}{\sum_p |p|^{-s}}\right] = \frac{1}{\Phi(m)} \left[1 - \frac{\sum_{p|m} |p|^{-s}}{\log(s-1) + O(1)} + \frac{\sum_{\chi \neq \chi_{\text{triv}}} \log L(s,\chi) + O(1)}{\log(s-1) + O(1)}\right]$.
  - Now, taking the limit as $s \to 1+$ makes the second term go to zero (since the numerator is finite) and the third term go to zero (since $L(1,\chi) \neq 0$ for $\chi \neq \chi_{\text{triv}}$), and so the value of the limit is just $1/\Phi(m)$, as claimed.

- We can, in fact, improve this argument to show that the natural density of the primes congruent to a modulo m is equal to $1/\Phi(m)$, not just the Dirichlet density.
  - To do this requires showing that $L(s,\chi)$ is zero-free on a larger region: specifically, we need it to be zero-free for $\text{Re}(s) = 1$, rather than just $s = 1$.
  - The L-function is in fact zero-free on a much larger region: as we will eventually prove, the only zeroes of $L(s,\chi)$ are on the line $\text{Re}(s) = 1/2$; this is the Riemann hypothesis for function fields.
  - Taking this zero-free result for granted, we again need to manipulate the series expressions for the $L(s,\chi)$. This time, we will use in a more substantial way the fact that the $L(s,\chi)$ for $\chi \neq \chi_{\text{triv}}$ are polynomials in $u = q^{-s}$ and compare the Euler products with their factorizations.

- Theorem (Strengthened Dirichlet Analogue): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the number of primes congruent to $a \pmod m$ having degree N is equal to $\frac{1}{\Phi(m)} \frac{q^N}{N} + O(\frac{q^{N/2}}{N})$, where the implied constant is independent of q and N.
  - If we only know that the L-function is zero free for $\text{Re}(s) > \theta$ for some $\theta \in (1/2, 1)$, we instead get an error term of $O(\frac{q^{\theta N}}{N})$, which is still good enough to establish that the natural density of primes congruent to $a \pmod m$ equals $1/\Phi(m)$.
  - Proof: For convenience, we first note the identity (*) $u \frac{\partial}{\partial u} \log(1 - \alpha u^d)^{-1} = \sum_{N=1}^{\infty} d\alpha^k u^{dN}$.
  - As we showed previously, if $\chi \neq \chi_{\text{triv}}$ then $L^*(u,\chi) = L(q^{-s},\chi)$ is a polynomial in $u = q^{-s}$ of degree at most $m-1$. Since its constant term is 1, we obtain a factorization of the form $L^*(u,\chi) = \prod_{i=1}^{m-1}(1 - \alpha_i(\chi)u)$ for some constants $\alpha_i(\chi) \in \mathbb{C}$.

- From the Euler product, we also have $L^*(u,\chi) = \prod_{p\nmid m}(1-\chi(p)u^{\deg p})^{-1} = \prod_{d=1}^{\infty}\prod_{p\nmid m,\deg p=d}(1-\chi(p)u^d)^{-1}$.
- Now apply the operator $u\frac{\partial}{\partial u}\log$ to the equality $\prod_{i=1}^{m-1}(1-\alpha_i(\chi)u) = \prod_{d=1}^{\infty}\prod_{p\nmid m,\deg p=d}(1-\chi(p)u^d)^{-1}$ and compare coefficients of u on both sides.
- For the LHS, using the identity (*) with $d=1$ yields $u\frac{\partial}{\partial u}\log L^*(u,\chi) = -\sum_{i=1}^{m-1}\sum_{N=1}^{\infty}\alpha_i(\chi)^N u^N = -\sum_{N=1}^{\infty}\left[\sum_{i=1}^{m-1}\alpha_i(\chi)^N\right]u^N$.
- Letting $c_N(\chi) = -\sum_{i=1}^{m-1}\alpha_i(\chi)^N$ yields the expansion $u\frac{\partial}{\partial u}\log\prod_{i=1}^{m-1}(1-\alpha_i(\chi)u) = \sum_{N=1}^{\infty}c_N(\chi)u^N$. For $\chi=\chi_{\text{triv}}$, we have $c_N(\chi) = q^N + O(1)$, while for $\chi\neq\chi_{\text{triv}}$, by the Riemann hypothesis we have $|\alpha_i(\chi)|\in\{q^0, q^{1/2}\}$ for each i, and so $c_N(\chi) = O(q^{N/2})$.
- For the RHS, we have

$$\begin{aligned} u\tfrac{\partial}{\partial u}\log L^*(u,\chi) &= \sum_{d=1}^{\infty}\sum_{p\nmid m,\deg p=d} u\tfrac{\partial}{\partial u}\log(1-\chi(p)u^d)^{-1} \\ &= \sum_{d=1}^{\infty}\sum_{p\nmid m,\deg p=d}\sum_{k=1}^{\infty} d\,\chi(p)^k u^{kd} \\ &= \sum_{N=1}^{\infty}\left[\sum_{d\mid N}\sum_{\deg p=N/d} d\,\chi(p)^{N/d}\right]u^N \end{aligned}$$

  by applying the identity (*) and then grouping together all of the terms of the same degree. This means $c_N(\chi) = \sum_{d\mid N}\sum_{\deg p=d} d\,\chi(p)^d$.
- Now, by separating out the terms with $d=1$ from the others, we see $c_N(\chi) = \sum_{d\mid N}\sum_{\deg p=N/d} d\,\chi(p)^{N/d} = N\sum_{\deg p=N}\chi(p) + \sum_{d\mid N,d\geq 2}\sum_{\deg p=N/d} d\,\chi(p)^d$. The absolute value of the second term is at most $\sum_{d\mid N,d\geq 2}\sum_{\deg p=N/d} d \leq \sum_{d\mid N,d\geq 2}\frac{q^{N/d}}{N/d} = O(q^{N/2})$.
- Therefore, we see $c_N(\chi) = N\sum_{\deg p=N}\chi(p) + O(q^{N/2})$.
- Now we use our Fourier decomposition from earlier: we have $\frac{1}{\Phi(m)}\sum_{\chi\in\hat{G}}\overline{\chi(a)}c_N(\chi) = N\cdot\#\{\text{primes } p\equiv a \pmod m\} + O(q^{N/2})$ using the expression we just computed.
- Also, we have $\sum_{\chi\in\hat{G}}\overline{\chi(a)}c_N(\chi) = q^N + O(q^{N/2})$ by directly summing over characters: $\chi=\chi_{\text{triv}}$ contributes the q^N term and the other characters each contribute $O(q^{N/2})$.
- Setting these two equal to one another yields $\#\{\text{primes } p\equiv a \pmod m\} = \frac{1}{\Phi(m)}\cdot\frac{q^N}{N} + O(\frac{q^{N/2}}{N})$, as claimed.

- Exercise: For $a, m\in\mathbb{F}_q[t]$ with a relatively prime to m, show that the proportion of primes of degree N congruent to $a \pmod m$ is $\frac{1}{\Phi(m)} + O(q^{-N/2})$, where the implied constant is independent of q and N.

0.7 (Sep 24) dth Powers and dth-Power Residue Symbols

- Our next task is to discuss the analogue of another famous result from elementary number theory: Gauss's celebrated law of quadratic reciprocity, along with its higher-order generalizations. A brief recap of the story over $\mathbb{Z}$:
  - If $a\in(\mathbb{Z}/p\mathbb{Z})^*$, we say a is a quadratic residue if $a\equiv b^2 \pmod p$ for some b, and otherwise we say a is a quadratic nonresidue.
  - Since the quadratic residues are simply the image of the squaring map on $(\mathbb{Z}/p\mathbb{Z})^*$, by the first isomorphism theorem there are $(p-1)/2$ of them. (One may also simply enumerate them as $1^2, 2^2, \ldots, [(p-1)/2]^2$.)

- The Legendre symbol $\left(\frac{a}{p}\right)$ is defined to be $+1$ on quadratic residues and -1 on quadratic nonresidues. By writing a as a power of the generator of $(\mathbb{Z}/p\mathbb{Z})^*$, one then obtains Euler's criterion: $a^{(p-1)/2} \equiv \left(\frac{a}{p}\right) \pmod{p}$, from which one sees that the Legendre symbol is multiplicative. Equivalently, it is a group homomorphism from $(\mathbb{Z}/p\mathbb{Z})^*$ to $\{\pm 1\}$.
- Exercise: Another group homomorphism from $(\mathbb{Z}/p\mathbb{Z})^*$ to $\{\pm 1\}$ is obtained by calculating the signature of the permutation associated to multiplication by a, as an element of the symmetric group S_{p-1}. Prove Zolotarev's lemma: this homomorphism is the same as the Legendre symbol.

- The law of quadratic reciprocity gives an unexpected relation between the Legendre symbols $\left(\frac{p}{q}\right)$ and $\left(\frac{q}{p}\right)$ for distinct odd primes p and q.
  - Explicitly, as first proven by Gauss, we have $\left(\frac{p}{q}\right)\left(\frac{q}{p}\right) = (-1)^{(p-1)(q-1)/4}$. Equivalently, $\left(\frac{p}{q}\right) = \left(\frac{q}{p}\right)$ if p or q is 1 mod 4, and otherwise $\left(\frac{p}{q}\right) = -\left(\frac{q}{p}\right)$ if both p, q are 3 mod 4.
  - A priori, it would seem that there is no reason for the values of $\left(\frac{p}{q}\right)$ and $\left(\frac{q}{p}\right)$ to be related to one another, since they are discussing seemingly independent questions (whether p is a square mod q and whether q is a square mod p).
  - But in fact, these questions are related: for $p^* = (-1)^{(p-1)/2}$, the value of $\left(\frac{p}{q}\right)$ determines whether the ideal (p) splits in the ring of integers $\mathcal{O}_{\sqrt{q^*}}$ of the quadratic extension $\mathbb{Q}(\sqrt{q^*})$ while the value of $\left(\frac{q}{p}\right)$ determines whether the ideal (q) splits in the ring of integers of the quadratic extension $\mathbb{Q}(\sqrt{p^*})$.
  - These two questions are related because there are several ways to understand the splitting of (q) in $\mathcal{O}_{\sqrt{p^*}}$.
  - First, from basic algebraic number theory, to determine whether (q) splits in $\mathcal{O}_{\sqrt{p^*}}$, one can study the splitting of the minimal polynomial $x^2 - x + \frac{1-p^*}{2}$ modulo q, which splits precisely when its discriminant p^* is a square: in other words, when $\left(\frac{p^*}{q}\right) = 1$.
  - Alternatively, one may look at the action of the local qth-power Frobenius map inside the Galois group of the cyclotomic field $\mathbb{Q}(\zeta_p)$, whose unique quadratic subfield is $\mathbb{Q}(\sqrt{p^*})$. Since the Galois group is cyclic, the Frobenius element Frob_q fixes $\mathbb{Q}(\sqrt{p^*})$ if and only if $q \in (\mathbb{Z}/p\mathbb{Z})^\times$ lies in $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/K)$. But this group is the unique index-2 subgroup of $(\mathbb{Z}/p\mathbb{Z})^*$, which is simply the quadratic residues, so this means (q) splits precisely when $\left(\frac{q}{p}\right) = 1$.
  - Comparing these two statements yields that $\left(\frac{p^*}{q}\right) = 1$ if and only if $\left(\frac{q}{p}\right) = 1$, and this can be shown to be equivalent to the usual version of quadratic reciprocity.
  - Exercise: For distinct odd primes p, q, show that $\left(\frac{p^*}{q}\right) = \left(\frac{q}{p}\right)$ is equivalent to $\left(\frac{p}{q}\right)\left(\frac{q}{p}\right) = (-1)^{(p-1)(q-1)/4}$, where $p^* = (-1)^{(p-1)/2}$.
  - There are very many other proofs of quadratic reciprocity, many of which involve lengthy formal manipulations of various sums and (generally) yield little to no intuition about why the result is actually true. There is a fairly nice proof using Gauss sums that, suitably interpreted, is really the same as the one given above.

- We would like to generalize the reciprocity law to handle general dth powers in $\mathbb{F}_q[t]$. We begin by describing the dth powers:

- Definition: If $f \in \mathbb{F}_q[t]$ is nonconstant and a is relatively prime to f, we say that a is a dth-power residue modulo f when $x^d \equiv a \pmod{f}$ has a solution for x. (In other words, when a is the dth power of something mod f.)

- Example: Over $\mathbb{F}_2[t]$, we see $t+1$ is a quadratic residue modulo t^3+t+1 since $t+1 \equiv (t^2+t+1)^2 \pmod{t^3+t+1}$.
- Example: Over $\mathbb{F}_5[t]$, we see $3t^2+3t+4$ is a cubic residue modulo t^3+t+1 since $3t^2+3t+4 \equiv (t^2+2t)^3 \pmod{t^3+t+1}$.
- By the Chinese remainder theorem, $x^d \equiv a \pmod{f}$ has a solution if and only if $x^d \equiv a \pmod{p^d}$ has a solution for each prime power p^d in the factorization of f.
- Thus, we need only consider the case where the modulus is a prime power, and we can handle this case fairly easily using our earlier analysis of the structure of $(A/p^dA)^*$.

- We can start by looking at the prime-modulus case, since it is the simplest.
  - As we have mentioned previously, $(A/pA)^*$ is the multiplicative group of the finite field A/pA, so this group has order $q^{\deg p}-1 = \tilde{p}^{f \deg p}-1$.
  - If d does not divide $|p|-1$, then the dth power map on $(A/pA)^*$ is injective by Lagrange's theorem, so it is a bijection, and so everything in $(A/pA)^*$ is a dth power.
  - This means we can ignore divisors of d that aren't factors of $|p|-1$, and so essentially we are reduced to the situation where d divides $|p|-1$.
  - By analogy with Euler's criterion in $\mathbb{Z}$, we would expect that the value of $a^{(|p|-1)/d}$ will identify whether or not a is a dth power. This is indeed the case:

- Proposition (dth Roots Mod p): If $p \in \mathbb{F}_q[t]$ is irreducible, a is not divisible by p, and d is a divisor of $|p|-1$, then $x^d \equiv a \pmod{p}$ is solvable if and only if $a^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - Proof 1: First, if $x^d \equiv a \pmod{p}$ then $a^{(|p|-1)/d} \equiv x^{|p|-1} \equiv 1 \pmod{p}$ by Euler.
  - For the converse, recall that we showed previously that $x^d \equiv 1 \pmod{p}$ has d solutions mod p whenever d divides $|p|-1$.
  - Therefore, the kernel of the dth-power map on $(A/pA)^*$ has size d, so by the first isomorphism theorem, the image, which is precisely the set of dth powers, has size $(|p|-1)/d$.
  - But by the same observation, there are exactly $(|p|-1)/d$ solutions to the equation $x^{(|p|-1)/d} \equiv 1 \pmod{p}$, so by the above, these must be exactly the dth powers.
  - Proof 2: As shown previously, $(A/pA)^*$ is cyclic of order $|p|-1$. Let u be a generator.
  - Since every element in $(A/pA)^*$ is a power of u, it is easy to see that for any d dividing $|p|-1$, the dth powers in $(A/pA)^*$ are precisely $\{u^d, u^{2d}, u^{3d}, \ldots, u^{d(|p|-1)d} = 1\}$. All of these elements clearly satisfy $x^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - Conversely, if $a = u^k$ has $a^{(|p|-1)/d} \equiv 1 \pmod{p}$, then $u^{k(|p|-1)/d} \equiv 1 \pmod{p}$ so since u has order $|p|-1$, d must divide k.

- Now that we have analyzed the prime case, the prime-power case follows by "lifting" the solutions from the prime case.
  - This is a consequence of a much more general result known as Hensel's lemma, which we might as well do in general.

- Proposition (Hensel's Lemma): If $p \in \mathbb{F}_q[t]$ is irreducible, $a \in \mathbb{F}_q[t]$, and $r(x)$ is any polynomial such that $r(a) \equiv 0 \pmod{p^d}$ and $r'(a) \not\equiv 0 \pmod{p}$, then there is a unique k modulo p such that $r(a+kp^d) \equiv 0 \pmod{p^{d+1}}$. Explicitly, if $u = f'(a)^{-1} \pmod{p}$, then $k = -\dfrac{uf(a)}{p^d}$.
  - By repeatedly applying Hensel's lemma, we can lift a solution of $r(a) \equiv 0 \pmod{p}$ to a solution modulo p^2, and then lift that to a solution modulo p^3, and so on and so forth, until we have a solution to the equation modulo any power of p.
  - This iteration process yields a sequence of solutions $x \equiv a_j \pmod{p^j}$ for each j, where $a_{j+1} = a_j - \dfrac{1}{r'(a)} r(a_j)$, which one may recognize as the iteration procedure from Newton's root-finding method. In fact, if we instead think of solving the polynomial $r(x) = 0$ p-adically (which amounts to taking the inverse limit $\varprojlim (A/p^dA)$), this lifting procedure is precisely Newton's method with starting point $x = a$.

- Proof: First, by the binomial theorem we have $(a+p^dk)^n = a^n + na^{n-1}p^dk + [\text{terms divisible by } p^{2d}] \equiv a^n + na^{n-1}p^dk \pmod{p^{d+1}}$.
- Then if $r(t) = \sum c_n t^n$ we see that $r(a+p^dk) \equiv \sum c_n(a^n + na^{n-1}p^dk) \equiv \sum c_n a^n + p^dk \sum nc_n a^{n-1} \equiv r(a) + p^dk \cdot r'(a) \pmod{p^{d+1}}$.
- By hypothesis, $r(a)+p^dk \cdot r'(a)$ is divisible by p^d. So dividing the congruence $r(a+kp^d) \equiv 0 \pmod{p^{d+1}}$ by p^d yields $\dfrac{r(a)}{p^d} + kr'(a) \equiv 0 \pmod{p}$, which has the unique solution $k \equiv -\dfrac{uf(a)}{p^d} \pmod{p}$, as claimed.

- This version of Hensel's lemma is quite a bit more than we really need here, but it will be helpful to have it available later.

- Corollary (dth Roots Mod p^e): If $p \in \mathbb{F}_q[t]$ is irreducible, d divides $|p|-1$, and p does not divide a, then $x^d \equiv a \pmod{p}$ has a root if and only if $x^d \equiv a \pmod{p^e}$ has a root for every $e \geq 1$.
  - Proof: If there is a solution to $x^d \equiv a \pmod{p^e}$ then clearly there is a solution mod p.
  - Conversely, if there is a solution mod p, then we claim we may lift the solution mod p^e using Hensel's lemma.
  - We just need to check that the derivative is not zero: for $r(x) = x^d$ we have $r'(a) = da^{d-1}$. Then $d \neq 0$ mod p because d divides $|p|-1 = \tilde{p}^{f \deg p} - 1$ and so d cannot be divisible by the characteristic $\tilde{p}$, and also $a \neq 0$ mod p because p does not divide a. Thus, Hensel's lemma applies, and we are done.

- Corollary (Counting dth Powers): If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $|p|-1$, then there are $\Phi(p^e)/d$ total dth-power residues modulo p^e.
  - Proof 1: Count residue classes: as shown earlier there are $(|p|-1)/d = \Phi(p)/d$ total dth-power residue classes modulo p. By the corollary above, the dth-power residue classes modulo p^e are precisely those that reduce to a dth power modulo p. So the probability of selecting one is $\Phi(p)/(d\,|p|)$, and thus the total number is $|p|^e \cdot \Phi(p)/(d\,|p|) = \Phi(p^e)/d$.
  - Proof 2 (sketch): The dth-power homomorphism commutes with reduction modulo p. Then just count the sizes of the various kernels and images and use the first isomorphism theorem.
  - Exercise: Show that for any monic polynomial m, there are $\Phi(m)/d^{\lambda(m)}$ total dth powers modulo m, where $\lambda(m)$ is the number of distinct monic irreducible factors of m.

- Returning back to the prime case, in the particular case where d divides $q-1$, then the dth roots of unity in $(A/pA)^*$ actually lie inside $\mathbb{F}_q$, because $x^d = 1$ already has d solutions inside $\mathbb{F}_q$ (since $\mathbb{F}_q^*$ is cyclic of order $q-1$).
  - We have shown above that a is a dth power modulo p if and only if $a^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - We can use this as the basis for our definition of the dth-power residue symbol, in analogy with Euler's criterion over $\mathbb{Z}$.

- Definition: If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $q-1$, then we define the dth-power residue symbol $\left(\dfrac{a}{p}\right)_d$ to be the unique element of $\mathbb{F}_q$ congruent to $a^{(|p|-1)/d}$ modulo p.
  - Example: For $d=2$ over $\mathbb{F}_3[t]$, we calculate $\left(\dfrac{t}{t^2+t+2}\right)_2 \equiv t^4 \equiv 2 \pmod{t^2+t+2}$.
  - Example: For $d=3$ over $\mathbb{F}_7[t]$, we calculate $\left(\dfrac{t}{t^2+2t+2}\right)_7 \equiv t^{16} \equiv 4 \pmod{t^2+2t+2}$.
  - Example: For $d=3$ over $\mathbb{F}_7[t]$, we calculate $\left(\dfrac{t}{t^2+t+6}\right)_7 \equiv t^{16} \equiv 1 \pmod{t^2+t+6}$, which means t is a cube modulo t^2+t+6.

- Proposition (Properties of Residue Symbols): If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $q-1$, the following hold:

1. $\left(\dfrac{a}{p}\right)_d = 0$ if and only if p divides a.

2. If $a \equiv b \pmod{p}$ then $\left(\frac{a}{p}\right)_d = \left(\frac{b}{p}\right)_d$.
3. The residue symbol is multiplicative: for any a, b, $\left(\frac{ab}{p}\right)_d = \left(\frac{a}{p}\right)_d \left(\frac{b}{p}\right)_d$.
4. $\left(\frac{a}{p}\right)_d = 1$ if and only if a is a dth-power residue modulo p.
5. If ζ is any dth root of unity in $\mathbb{F}_q$, then there exists $a \in \mathbb{F}_q[t]$ with $\left(\frac{a}{p}\right)_d = \zeta$.
6. The residue symbol is a surjective group homomorphism from $(A/pA)^*$ to μ_d, the group of dth roots of unity in $\mathbb{F}_q$.
7. If $d|d'$ then $\left(\frac{a}{p}\right)_d = \left(\frac{a}{p}\right)_{d'}^{d'/d}$.
8. If $\alpha \in \mathbb{F}_q$ then $\left(\frac{\alpha}{p}\right)_d = \alpha^{(q-1)/d \cdot \deg p}$.

- Proofs: (1)-(4) are trivial from the definition or results previously shown. (5) follows by the first isomorphism theorem, since the kernel of the $(|p|-1)/d$th-power map has size $(|p|-1)/d$ hence the image has size d. (6) is a rephrasing of (3) and (5).
- (7) follows by noting $\left(\frac{a}{p}\right)_{d'}^{d'/d} \equiv (a^{(|p|-1)/d'})^{d'/d} = a^{(|p|-1)/d} \equiv \left(\frac{a}{p}\right)_d \pmod{p}$, and then observing that since the residue symbols are both elements of $\mathbb{F}_q$, the congruence mod p forces actual equality.
- For (8), first note that $\frac{|p|-1}{d} = \frac{q^{\deg p}-1}{d} = (1+q+q^2+\cdots+q^{\deg p-1})(q-1)/d$. Then since $\alpha^q = \alpha$ by Fermat's little theorem in $\mathbb{F}_q$, we have $\left(\frac{\alpha}{p}\right)_d \equiv \alpha^{(|p|-1)/d} = (\alpha \cdot \alpha^q \cdot \alpha^{q^2} \cdot \cdots \cdot \alpha^{q^{\deg p-1}})^{(q-1)/d} = \alpha^{\deg p \cdot (q-1)/d} \pmod{p}$. Then as in (7), the congruence modulo p forces equality.

- We can now state the dth-power reciprocity law, which we will prove next time:
- Theorem (dth-Power Reciprocity): If d divides $q-1$ and P, Q are monic irreducible polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{Q}{P}\right)_d = (-1)^{(q-1)(\deg P)(\deg Q)/d} \left(\frac{P}{Q}\right)_d$.

0.8 (Sep 29) The dth-Power Reciprocity Law

- To prove the reciprocity law, we first need a reciprocity result about roots of polynomials known as Weil reciprocity:
- Lemma (Weil Reciprocity): If $P(t) = (t-r_1)\cdots(t-r_n)$ and $Q(t) = (t-s_1)\cdots(t-s_n)$ are monic polynomials over a field F, with the $r_i, s_j \in F$, then $\prod_{i=1}^n Q(r_i) = (-1)^{(\deg P)(\deg Q)} \prod_{j=1}^m P(s_j)$.
  - Proof: Note that $Q(r_i) = \prod_{j=1}^m (r_i - s_j)$ so $\prod_{i=1}^n Q(r_i) = \prod_{i=1}^n \prod_{j=1}^m (r_i - s_j)$. In the same way, $\prod_{j=1}^m P(s_j) = \prod_{j=1}^m \prod_{i=1}^n (s_j - r_i)$.
  - These expressions are the same up to switching the order of the products and scaling each of the $mn = (\deg P)(\deg Q)$ terms by -1, so the result follows.
- We can now prove the dth-power reciprocity law:
- Theorem (dth-Power Reciprocity): If d divides $q-1$ and P, Q are monic irreducible polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{Q}{P}\right)_d = (-1)^{(q-1)(\deg P)(\deg Q)/d} \left(\frac{P}{Q}\right)_d$.
  - The main idea of the proof is to exploit properties of the Frobenius map on the roots of P and Q in their splitting field over $\mathbb{F}_q$, and then use Weil reciprocity.

- ◦ Proof: From property (7) of the residue symbol, we have $\left(\frac{a}{p}\right)_d = \left(\frac{a}{p}\right)_{d'}^{d'/d}$, so it is enough to prove the reciprocity law when $d = q-1$.
- ◦ Now let α be a root of P and β be a root of Q in a splitting field $E/\mathbb{F}_q$ for the polynomial PQ.
- ◦ Since $E/\mathbb{F}_q$ is a finite-degree extension of a finite field, its Galois group is cyclic and generated by the qth-power Frobenius map.
- ◦ Also, since P and Q are irreducible over $\mathbb{F}_q$, we must have the factorizations

$$P(t) = (t-\alpha)(t-\alpha^q)(t-\alpha^{q^2})\cdots(t-\alpha^{q^{\deg P-1}})$$
$$Q(t) = (t-\beta)(t-\beta^q)(t-\beta^{q^2})\cdots(t-\beta^{q^{\deg Q-1}})$$

since α, α^q, α^{q^2}, ... are all the Galois conjugates of α and P is irreducible (with the same logic applying to β and Q).

- ◦ Inside $E[t]$, we have $\left(\frac{Q}{P}\right)_{q-1} \equiv [Q(t)]^{(q^{\deg P}-1)/(q-1)} = [Q(t)]^{1+q+q^2+\cdots+q^{\deg P-1}} = Q(t)Q(t)^qQ(t)^{q^2}\cdots Q(t)^{q^{\deg P-1}} \equiv Q(t)Q(t^q)Q(t^{q^2})\cdots Q(t^{q^{\deg P-1}}) \pmod P$ since $Q(t^q) = Q(t)^q$ in characteristic q.
- ◦ Reducing both sides modulo the factor $t-\alpha$ of P (equivalently, evaluating both sides at $t=\alpha$) then yields $\left(\frac{Q}{P}\right)_{q-1} \equiv Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) \pmod{t-\alpha}$. Since the right-hand side of this expression is the product of the values of Q evaluated at the roots of P, it is the same for any other root of P we choose in place of α.
- ◦ So by the Chinese remainder theorem, in fact $\left(\frac{Q}{P}\right)_{q-1} \equiv Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) \pmod P$. But the right-hand side is an element of E, and since it is a $(q-1)$st root of unity (or alternatively, since it is Galois-invariant), it must actually be in $\mathbb{F}_q$. So since these quantities are congruent modulo P, they must actually be equal as elements of $\mathbb{F}_q$.
- ◦ This means $\left(\frac{Q}{P}\right)_{q-1} = Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1})$. In the same way, $\left(\frac{P}{Q}\right)_{q-1} = P(\beta)P(\beta^q)\cdots P(\beta^{q^{\deg Q}-1})$.
- ◦ Weil reciprocity then says $Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) = (-1)^{(\deg P)(\deg Q)}P(\beta)P(\beta^q)\cdots P(\beta^{q^{\deg Q}-1})$, so we see $\left(\frac{Q}{P}\right)_{q-1} = (-1)^{(\deg P)(\deg Q)}\left(\frac{P}{Q}\right)_{q-1}$, which establishes the case $d = q-1$.
- ◦ The case where d divides $q-1$ follows immediately and gives the general statement above.

- Just as in the case of $\mathbb{Q}$, to give a convenient method for calculating residue symbols, we can extend the definition to include nonprime moduli (i.e., generalizing the Jacobi symbol):
- Definition: If $b \in \mathbb{F}_q[t]$ has prime factorization $b = uq_1^{b_1}\cdots q_n^{b_n}$ for distinct monic irreducible q_i and $u \in \mathbb{F}_q^\times$, then we define the general residue symbol as $\left(\frac{a}{b}\right)_d = \prod_{j=1}^n \left(\frac{a}{q_i}\right)^{b_i}$.
- Proposition (Properties of Residue Symbols, II): If $b \in \mathbb{F}_q[t]$ is nonzero and d divides $q-1$, the following hold:

1. $\left(\frac{a}{b}\right)_d$ is either 0 or a dth root of unity, and $\left(\frac{a}{b}\right)_d \neq 0$ if and only if a, b are relatively prime.
2. If $a_1 \equiv a_2 \pmod b$ then $\left(\frac{a_1}{b}\right)_d = \left(\frac{a_2}{b}\right)_d$.
3. The residue symbol is multiplicative on the top: $\left(\frac{a_1a_2}{b}\right)_d = \left(\frac{a_1}{b}\right)_d\left(\frac{a_2}{b}\right)_d$.
4. The residue symbol is multiplicative on the bottom: $\left(\frac{a}{b_1b_2}\right)_d = \left(\frac{a}{b_1}\right)_d\left(\frac{a}{b_2}\right)_d$.
5. If $\gcd(a,b) = 1$ and a is a dth-power residue modulo b, then $\left(\frac{a}{b}\right)_d = 1$. (The converse need not hold.)

6. If $d|d'$ then $\left(\frac{a}{b}\right)_d = \left(\frac{a}{b}\right)_{d'}^{d'/d}$.
7. If $\alpha \in \mathbb{F}_q$ then $\left(\frac{\alpha}{b}\right)_d = \alpha^{(q-1)/d \cdot \deg b}$.

- Proofs: (1)-(4) follow straightforwardly from the definition, while (6) and (7) follow the same way as for the residue symbol with prime modulus. For (5), if $a \equiv c^d \pmod{p}$ then $\left(\frac{a}{b}\right)_d = \left(\frac{c^d}{b}\right)_d = \left(\frac{c}{b}\right)_d = 1$ since $\left(\frac{c}{b}\right)$ is a dth root of unity (since it is not zero since a, b are relatively prime).
- We will remark that the residue symbol $\left(\frac{\star}{b}\right)_d : (A/bA)^* \to \mu_d$ is still a group homomorphism since it is multiplicative by (3), but it is not necessarily surjective when b is not prime. For example, if $b = p^d$ is a dth power, then by (4) we see that $\left(\frac{a}{b}\right)_d = \left(\frac{a}{p}\right)_d^d = 1$ for all $a \in (A/bA)^*$. (This also shows that the converse of (5) is false, as noted above.)

- We can write down the reciprocity law for general dth-power residue symbols:
- Theorem (General Reciprocity Law): If d divides $q-1$ and a, b are any nonzero polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{a}{b}\right)_d = (-1)^{(q-1)(\deg a)(\deg b)/d} [\mathrm{sgn} a]^{(q-1)/d \cdot \deg b} [\mathrm{sgn} b]^{-(q-1)/d \cdot \deg a} \left(\frac{b}{a}\right)_d$.
  - Proof (sketch): As in the prime case, reduce to the case $d = q-1$. Then pull out the leading coefficients of a, b (these are where the $\mathrm{sgn} a$ and $\mathrm{sgn} b$ terms come from) and then apply the definition of the general residue symbol to write $\left(\frac{a}{b}\right)_{q-1}$ and $\left(\frac{b}{a}\right)_{q-1}$ as products of residue symbols with prime moduli, apply the prime-modulus reciprocity law, and tally up the results. The full details are left as an exercise.

- A standard application of quadratic reciprocity over $\mathbb{Z}$ is to characterize all of the prime moduli for which a given integer m is a quadratic residue.
  - Typical examples of such statements: -1 is a quadratic residue mod p when $p \equiv 1 \pmod 4$, 3 is a quadratic residue mod p when $p \equiv 1, 11 \pmod{12}$, 5 is a quadratic residue mod p when $p \equiv 1, 4 \pmod 5$, and so forth.
  - Aside from the special cases of -1 and 2, one may answer this question simply by factoring m as a product of primes $m = q_1 \cdots q_k$, so that $\left(\frac{m}{p}\right) = \left(\frac{q_1}{p}\right) \cdots \left(\frac{q_k}{p}\right)$, and then applying quadratic reciprocity to flip each of the quadratic residue symbols. The end result is that the statement $\left(\frac{m}{p}\right) = +1$ is equivalent to a congruence condition for p modulo $4m$, which one may calculate explicitly if desired.

- We can use this same type of argument to solve the analogous problem in function fields:
- Theorem (Criterion for dth-Power Residues): Let $m \in \mathbb{F}_q[t]$ be monic and $d|(q-1)$, and let $\{a_1, \ldots, a_k\}$ be coset representatives for the residue classes in $(A/mA)^*$ with $\left(\frac{a}{m}\right)_d = +1$ and $\{b_1, \ldots, b_k\}$ be coset representatives for the residue classes in $(A/mA)^*$ with $\left(\frac{b}{m}\right)_d = -1$ (if there are any). Then the following hold:

1. If $\deg(m)$, $(q-1)/d$, or $\mathrm{char}(\mathbb{F}_q)$ is even, then m is a dth power modulo an irreducible monic polynomial p if and only if $p \equiv a_i \pmod m$ for some i.
2. If $\deg(m)$, $(q-1)/d$, and $\mathrm{char}(\mathbb{F}_q)$ are all odd, then m is a dth power modulo an irreducible monic polynomial p if and only if either $\deg(p)$ is even and $p \equiv a_i \pmod m$ for some i, or $\deg(p)$ is odd and $p \equiv b_i \pmod m$ for some i.
   - Proof: Note that $p \equiv a_i \pmod m$ is equivalent to saying $\left(\frac{p}{m}\right)_d = 1$, while $p \equiv b_i \pmod m$ is equivalent to saying $\left(\frac{p}{m}\right)_d = -1$.

- Since p and m are monic, by the reciprocity law we see $\left(\frac{m}{p}\right)_d = (-1)^{(q-1)/d \cdot \deg(m)\deg(p)} \left(\frac{p}{m}\right)_d$.
- First, if q is even, then $\mathrm{char}(\mathbb{F}_q) = 2$: then $-1 = 1$ over $\mathbb{F}_q$, so $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$. Likewise, if $\deg(m)$ or $(q-1)/d$ is even, then the exponent of -1 is even, so again we see $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$. Together with the observation above, (1) follows.
- For (2), if $\deg(m)$, $(q-1)/d$, and $\mathrm{char}(\mathbb{F}_q)$ are all odd, then $-1 \neq 1$ and $(-1)^{(q-1)/d \cdot \deg(m)\deg(p)} = (-1)^{\deg p}$. So $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$ if $\deg(p)$ is even while $\left(\frac{m}{p}\right)_d = -\left(\frac{p}{m}\right)_d$ if $\deg(p)$ is odd. This yields (2).

- Example: Identify all monic irreducibles $p \in \mathbb{F}_3[t]$ such that t is a square modulo p.
  - There are two residue classes in $(A/tA)^*$, namely 1 and 2, and we see $\left(\frac{1}{t}\right)_2 = 1$ while $\left(\frac{2}{t}\right)_2 = -1$.
  - Since $\deg(m) = 1$, $(q-1)/d = 1$, and $\mathrm{char}(\mathbb{F}_q) = 3$, we are in case (2). Thus, m is a quadratic residue modulo the monic irreducible polynomial p precisely when $\deg(p)$ is odd and $p \equiv 2 \pmod{t}$, or when $\deg(p)$ is even and $p \equiv 1 \pmod{t}$.
  - For example, we see that t is a square modulo the irreducible polynomial $t^3 + 2t + 2 \in \mathbb{F}_3[t]$, and indeed with some more work, one may calculate $t \equiv (t^2 + t + 2)^2 \pmod{t^3 + 2t + 2}$.
  - Exercise: Extend this example to describe all monic irreducibles $p \in \mathbb{F}_q[t]$ such that t is a square modulo p for arbitrary finite fields $\mathbb{F}_q$.

- Another interesting application of the dth-power reciprocity law is to establish a "Hasse principle"-type result for dth powers.
  - Obviously, if a polynomial with integer coefficients has a solution in $\mathbb{Z}$, then it also has solutions modulo p^k for all prime powers p^k (equivalently, it has a p-adic solution for each p) and it also has a real solution.
  - The Hasse principle asks when the converse of this observation is valid: if a polynomial has a p-adic root and a real root, does it necessarily have a rational root? The general idea is that one may try to piece together information modulo the prime powers for many primes p using the Chinese remainder theorem, but it is not clear when this actually forces the existence of a global solution.
  - As first proven by Minkowski for integer coefficients (and then later extended by Hasse for number-field coefficients), for quadratic polynomials this local-global principle holds: if a quadratic polynomial has a p-adic root and a real root, it necessarily has a rational root.
  - The result is known to be false for cubic forms: Selmer's famous counterexample is the cubic equation $3x^3 + 4y^3 + 5z^3 = 0$, which has no rational solution but does have real solutions and p-adic solutions for all p.
  - Even in the absence of a literal Hasse-principle statement, in many cases one can analyze the precise obstructions to lifting local solutions to global solutions. (An example of this sort of obstruction can be found in the statement of the Grunwald-Wang theorem.)

- Theorem (Hasse Principle for dth Powers): Let $m \in \mathbb{F}_q[t]$ have positive degree and $d|(q-1)$. If $x^d \equiv m \pmod{p}$ is solvable for all but finitely many irreducible polynomials p, then $x^d = m$ has a solution in $\mathbb{F}_q[t]$ (i.e., m is globally a dth power).
  - Proof: Let $m = \beta q_1^{d_1} \cdots q_k^{d_k}$ where the q_i are distinct monic irreducibles and β is a constant. We first show that if any d_i is not divisible by d, then there are infinitely many irreducibles p such that $\left(\frac{m}{p}\right)_d \neq 1$.
  - To show this, suppose without loss of generality that d_1 is not divisible by d. We inductively construct an infinite set of irreducibles $\{r_i\}$ with $\left(\frac{m}{r_i}\right)_d \neq 1$, so suppose we have a set (possibly empty to start) $\{r_1, \ldots, r_s\}$ of monic irreducibles not dividing m with $\left(\frac{m}{r_i}\right)_d \neq 1$ for all i.

- Select any primitive dth root of unity ζ_d: then there exists an element $c \in \mathbb{F}_q[t]$ with $\left(\frac{c}{q_i}\right)_d = \zeta_d$ by our properties of the dth-power residue symbol.
- By the Chinese remainder theorem, there exist solutions a to the system of congruences $a \equiv c \pmod{q_1}$, $a \equiv 1 \pmod{q_2 \cdots q_k}$, $a \equiv 1 \pmod{r_1 \cdots r_s}$. Select any such solution that is monic and has degree divisible by $2d$.
- For this a, we have $\left(\frac{a}{m}\right)_d = \prod_{i=1}^{k} \left(\frac{a}{q_i}\right)_d^{d_i} = \zeta_d^{d_1} \neq 1$ since d_1 is not divisible by d.
- Then by the reciprocity law, we then have $\left(\frac{m}{a}\right)_d = (-1)^{(q-1)/d \cdot (\deg m)(\deg a)} \left(\frac{a}{m}\right)_d = \left(\frac{a}{m}\right)_d \neq 1$, since the exponent of -1 has a factor of 2 from $\deg a$.
- Since the general dth-power residue symbol is multiplicative on the bottom, there must be some monic irreducible factor r_{s+1} of a such that $\left(\frac{m}{r_{s+1}}\right) \neq 1$ since $\left(\frac{a}{m}\right)_d \neq 1$. This monic irreducible factor is relatively prime to $r_1 \cdots r_s$ since $a \equiv 1 \pmod{r_1 \cdots r_s}$, so we have found another monic irreducible to add to our list.
- By induction, we can construct infinitely many such irreducibles.
- Now, if $x^d \equiv m \pmod{p}$ is solvable for all but finitely many irreducible polynomials p, then by the above, each of the exponents d_i must be divisible by d. This means $m = \beta \cdot \tilde{m}^d$ for some monic polynomial $\tilde{m}$, so all that remains is to show that β is a dth power.
- For any irreducible p not dividing m, we have $\left(\frac{m}{p}\right)_d = \left(\frac{\beta}{p}\right)_d = \beta^{(q-1)/d \cdot \deg p}$ as we have previously shown. Since there are irreducibles of any desired degree in $\mathbb{F}_q[t]$, select p to be one of degree relatively prime to d with $\left(\frac{m}{p}\right)_d = 1$: then $\beta^{(q-1)/d \cdot \deg p} = 1$ implies $\beta^{(q-1)/d} = 1$, which is equivalent to saying that β is a dth power. Then m itself is a dth power, as claimed.

0.9 (Oct 1) Transcendence and Localization

- We now move into the second major part of the course, which deals with algebraic function fields: these are function fields of transcendence degree 1 over a general constant field F.
  - Later, we will specialize to function fields over $\mathbb{F}_q$ (equivalently, these are the finite-degree field extensions of $\mathbb{F}_q(t)$), which along with algebraic number fields (the finite-degree field extensions of $\mathbb{Q}$) constitute the global fields.
  - Global fields (to be considered as parallel to local fields) share a number of common properties that we will elucidate and study.
- We begin by reviewing some basic facts about transcendental extensions.
- Definition: Let K/F be a field extension. We say a subset S of K is algebraically dependent over F if there exists a finite subset $\{s_1, \ldots, s_n\} \in S$ and a nonzero polynomial $p \in F[x_1, \ldots, x_n]$ such that $p(s_1, \ldots, s_n) = 0$. If there exists no such p for any finite subset of S, we say S is algebraically independent.
  - The general idea here is that a set of elements is algebraically dependent if they satisfy some algebraic (i.e., polynomial) relation over F.
  - Example: If $x_1, \ldots, x_n$ are indeterminates inside $F(x_1, \ldots, x_n)$, the function field in n variables, then the set $\{x_1, \ldots, x_n\}$ is algebraically independent over F.
  - Example: Over $\mathbb{Q}$, the set $\{\pi, \pi^2\}$ is algebraically dependent, since $p(x, y) = x^2 - y$ has $p(\pi, \pi^2) = 0$.
  - Example: Over $\mathbb{Q}$, the set $\{\sqrt[3]{2}\}$ is algebraically dependent, since $p(x) = x^3 - 2$ has $p(\sqrt[3]{2}) = 0$.
  - More generally, the set $\{\alpha\}$ is algebraically independent over F if and only if α is transcendental over F.
  - Exercise: Show that the set $\{x + y, x^2 + y^2\}$ is algebraically independent in $F(x, y)$ for any field F of characteristic not 2, but is algebraically dependent if F has characteristic 2.

- Example: In $F(x,y)$, the set $\{x+y, x^2+y^2, x^3+y^3\}$ is algebraically dependent, since $p(a,b,c) = a^3 - 3ab + 2c$ has $p(x+y, x^2+y^2, x^3+y^3) = 0$.

- The notion of algebraic independence generalizes the notion of linear independence, and as such the two concepts are related in various ways.
  - It is easy to see that any subset of an algebraically independent set is algebraically independent, while any set containing an algebraically dependent set is algebraically dependent.
  - Since having a basis of a vector space is very convenient for calculations, we might therefore hope to define an analogous "transcendence basis" to be an algebraically independent set that generates the extension K/F.
  - Unfortunately, such a set need not exist: for example, $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$ has no such set, because there are no transcendental elements at all.
  - The correct analogy is instead to observe that a basis for a vector space is a maximal linearly independent set:

- Definition: Let K/F be a field extension. A transcendence base for K/F is an algebraically independent subset S of K that is maximal in the set of all algebraically independent subsets of K.
  - Remark: The term "transcendence basis" is also used occasionally.
  - By a straightforward Zorn's lemma argument, every extension has a transcendence base. (Exercise: Write down this argument.)
  - Example: The empty set $\emptyset$ is a transcendence base for $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$. More generally, K/F is algebraic if and only if $\emptyset$ is a transcendence base.
  - Example: The set $\{x\}$ is a transcendence base for $F(x)$ over F.

- Here are some of the fundamental properties of transcendence bases, many of which are analogous to properties of vector spaces:

- Proposition (Transcendence Bases): Suppose K/F is a field extension and S is a subset of K.

1. If S is algebraically independent and $\alpha \in K$, then $S \cup \{\alpha\}$ is algebraically independent over F if and only if α is transcendental over $F(S)$.
   - This is the algebraic analogue of the statement that if S is linearly independent, then $S \cup \{\alpha\}$ is linearly independent if and only if α is not in the span of S.
   - Proof: Suppose $S \cup \{\alpha\}$ is algebraically dependent. Then there exists $s_i \in S$ and $p \in F[x]$ with $p(\alpha, s_1, \ldots, s_n) = 0$ and $p \neq 0$. View p as a polynomial in its first variable with coefficients in $F[s_1, \ldots, s_n]$: there must be at least one term involving α, as otherwise p would give an algebraic dependence in S. Then α is the root of a nonzero polynomial with coefficients in $F[s_1, \ldots, s_n] \subseteq F(s_1, \ldots, s_n) \subseteq F(S)$, so it is algebraic over $F(S)$.
   - Conversely, suppose that α is algebraic over $F(S)$. Then α is the root of some nonzero polynomial with coefficients in $F(S)$. Each coefficient of this polynomial is an element of $F(S)$; clearing denominators yields a nonzero polynomial p with coefficients in $F[s_1, \ldots, s_n]$ for the elements $s_i \in S$ that appear in these coefficients. This polynomial yields an algebraic dependence in $S \cup \{\alpha\}$.
2. S is a transcendence base of K/F if and only if K is algebraic over $F(S)$.
   - Proof: This follows from (1) and the maximality of transcendence bases: S is a transcendence base if and only if no elements in K can be adjoined to S while preserving algebraic independence, and by (1) this is equivalent to saying that all elements in K are algebraic over $F(S)$.
3. If T is a subset of K such that $K/F(T)$ is algebraic, then T contains a transcendence base of K/F.
   - Proof: Apply Zorn's lemma to the collection of all algebraically independent subsets of T, partially ordered by inclusion.
   - A maximal element M in this collection must then be a transcendence base for K/F: if $\beta \in K$ then β must be algebraic over $K/F(M)$ by the maximality of M, and then M is a transcendence base by (2).

4. If T is an algebraically independent subset of K, then T can be extended to a transcendence base of K/F.
    - Proof: This is the analogue of the fact that every linearly independent subset can be extended to a basis, and the proof follows from a similar Zorn's lemma argument.
5. If $S = \{s_1, \ldots, s_n\}$ is a transcendence base for K/F and $T = \{t_1, \ldots, t_m\}$ is any algebraically independent set, then there is a reordering of S, say $\{a_1, \ldots, a_n\}$, such that for each $1 \leq k \leq m$, the set $\{t_1, t_2, \ldots, t_k, a_{k+1}, \ldots, a_n\}$ is a transcendence base for K/F.
    - Proof: This is the analogue of the replacement theorem for linearly independent sets, and the proof proceeds inductively in essentially the same way. (We will omit the details.)
6. If S is a (finite) transcendence base for K/F, then any subset T of K having larger cardinality than S must be algebraically dependent.
    - Proof: If $S = \{s_1, \ldots, s_n\}$ is finite, apply the replacement theorem (5) to S and T. At the end of the replacement, the result is that $\{t_1, \ldots, t_n\}$ is a transcendence base. But then by (2), any additional element of T would be algebraic over $\{t_1, \ldots, t_n\}$, contradicting the algebraic independence of T.
7. Any two transcendence bases S and T for K/F have the same cardinality.
    - Proof: If the bases are infinite the result is immediate. If S has finite cardinality n, then the result follows by applying (6), since then T's cardinality m must satisfy $m \leq n$ (since T is algebraically independent and S is a transcendence base) and also $n \leq m$ (since S is algebraically independent and T is a transcendence base).

- The result of the last part of the proposition shows that any two transcendence bases have the same cardinality, and in analogy with the situation for vector spaces, this cardinality behaves somewhat like an extension degree:
- Definition: Let K/F be a field extension. The transcendence degree of K/F, denoted $\text{trdeg}(K/F)$, is the cardinality of any transcendence base of K/F.
- The key property of transcendence degree is that it is additive in towers:
- Proposition (Transcendence in Towers): If $L/K/F$ is a tower of extensions, then $\text{trdeg}(L/F) = \text{trdeg}(L/K) + \text{trdeg}(K/F)$.
    - The idea here is quite simple: we want to show that the union of transcendence bases for K/F and L/K gives a transcendence base for L/F.
    - Proof: First suppose that both $\text{trdeg}(K/F)$ and $\text{trdeg}(L/K)$ are finite, and let $S = \{s_1, \ldots, s_n\}$ and $T = \{t_1, \ldots, t_m\}$ be transcendence bases for K/F and L/K. Then $S \cap T = \emptyset$ since each t_i is transcendental over K.
    - Furthermore, K is algebraic over $F(S)$, so $K(T)$ is algebraic over $F(T)(S) = F(S \cup T)$ by our results on algebraic extensions.
    - Then since L is algebraic over $K(T)$, we deduce that L is algebraic over $F(S \cup T)$, also by our results on algebraic extensions.
    - Thus, by property (3) above, $S \cup T$ contains a transcendence base of L/F.
    - Finally, we claim $S \cup T$ is algebraically independent over F, so suppose that $p(s_1, \ldots, s_n, t_1, \ldots, t_m) = 0$ for some $p \in F[x_1, \ldots, x_n, y_1, \ldots, y_m]$.
    - Separate monomial terms to write $p(s_1, \ldots, s_n, t_1, \ldots, t_m) = 0$ as a sum $\sum f_i(s_1, \ldots, s_n) g_i(t_1, \ldots, t_m) = 0$ with $f_i \in F[x_1, \ldots, x_n]$ and $g_i \in F[y_1, \ldots, y_m]$.
    - Now, since T is algebraically independent over $F(S) \subseteq K$, all of the $f_i(s_1, \ldots, s_n)$ must be zero (as elements of K). But since S is algebraically independent over F, that means all of the polynomials $f_i(x_1, \ldots, x_n)$ must be zero (as polynomials).
    - This means p is the zero polynomial, and so $S \cup T$ is algebraically independent.

- Fields that are generated by a transcendence base are particularly convenient:
- Definition: The extension K/F is purely transcendental if $K = F(S)$ for some transcendence base S of K/F.

- Equivalently, K/F is purely transcendental when it is generated (as a field extension) by an algebraically independent set.
- If $S = \{s_1, \ldots, s_n\}$, then the purely transcendental extension $K = F(S)$ is ring-isomorphic to the function field $F(x_1, \ldots, x_n)$ in n variables: it is not hard to check that the map sending s_i to x_i is an isomorphism.
- If K/F has transcendence degree 1 or 2 and E/F is an intermediate extension, then in fact E is also purely transcendental: the degree-1 case is a theorem of Lüroth that we will prove later, while the degree-2 case is a theorem of Castelnuovo. In higher degrees, there do exist extensions that are not purely transcendental, but it is not easy to verify this fact.

- Now let F be a field and K be an extension of F of transcendence degree 1.
  - By the results above, there exists $x \in K$ such that $K/F(x)$ has transcendence degree 0, which is to say, it is algebraic.
  - Since we do not want to worry for the moment about infinite-degree algebraic extensions, we will make the further assumption that this extension $K/F(x)$ has finite degree.

- Definition: We say K is an (algebraic) function field over F if there exists $x \in K$ such that x is transcendental over F and $K/F(x)$ is finite.
  - Example: $\mathbb{Q}(x)$ is an algebraic function field over $\mathbb{Q}$.
  - Example: $\mathbb{C}(x, \sqrt{x^2 - 1})$ is an algebraic function field over $\mathbb{C}$.
  - Note that the algebraic closure of F inside K has finite degree over F: this follows by noting that if E/F is algebraic inside K, then $[E : F] = [E(x) : F(x)] \leq [K : F(x)] < \infty$.
  - So, without loss of generality, we may replace F by its algebraic closure inside K. In this case we call F the constant field of K.
  - If F is the constant field of K, then there are no elements of K that are algebraic over F other than the elements of F themselves. Equivalently, every element of $K \backslash F$ is transcendental over F.
  - Finally, since the transcendence degree of K/F is 1, for any two $a, b \in K \backslash F$, there is some nonzero polynomial $g \in F[x, y]$ such that $F[a, b] = 0$.

0.10 (Oct 6) Localization, Discrete Valuations

- Now that we have some very basic facts about function fields, our goal is to do number theory.
  - In order to do this, however, we need to know how to define primes in the function field context.
  - Over $\mathbb{Q}$, the primes arise as the prime ideals of the ring of integers $\mathbb{Z}$, which we can define starting from $\mathbb{Q}$ purely in terms of integral closures. For other number fields, we also define their primes using integral closures.
  - However, this approach will not work for function fields, because (as noted above) everything in K not in F is transcendental over F, so there is no sensible way to define a "ring of integers" inside K using integrality.
  - Instead, we need to use a different sort of construction to give a sensible notion of a prime: that of a discrete valuation on K.

- In order to develop all of this properly, we first need to review some facts about localization.

- Proposition (Localization): Let R be a commutative ring with 1 and D be a multiplicatively closed subset of R containing 1. Then there exists a commutative ring $D^{-1}R$, the localization of R at D, and a ring homomorphism $\pi : R \to D^{-1}R$ such that any for any ring homomorphism $\psi : R \to S$ sending 1 to 1 and such that $\psi(d)$ is a unit in S for every $d \in D$, there exists a unique homomorphism $\Psi : D^{-1}R \to S$ such that $\Psi \circ \pi = \psi$.
  - More succinctly, any homomorphism $\psi : R \to S$ such that ψ maps all of the elements of D into units necessarily extends to a homomorphism $\Psi : D^{-1}R \to S$.

- The main idea is simply to define "fractions" r/d with $r \in R$ and $d \in D$ via an appropriate equivalence relation, and then to write down the usual rules of fraction arithmetic and verify that all of the definitions are well posed.
- Proof (outline): Define an equivalence relation on elements of $R \times D$ by setting $(r,d) \sim (s,e)$ whenever there exists $y \in D$ such that $y(ds - er) = 0$; it is straightforward to check that $\sim$ is an equivalence relation.
- Denote the equivalence class of (r,d) by the symbol r/d and the set of all equivalence classes by $D^{-1}R$, and define the two operations $r/d + s/e = (re+ds)/(de)$ and $r/d \cdot s/e = (rs)/(de)$ on $D^{-1}R$. It is tedious but straightforward to see that these operations make $D^{-1}R$ into a commutative ring with 1.
- Now define $\pi(r) = r/1$ and suppose $\Psi : D^{-1}R \to S$ is a homomorphism with $\Psi \circ \pi = \psi$.
- Then we must have $\Psi(r/1) = (\Psi \circ \pi)(r) = \psi(r)$, and also $1 = \Psi(1/1) = \Psi(1/d)\Psi(d/1)$, meaning that $\Psi(1/d) = \psi(d)^{-1}$. Then $\Psi(r/d) = \Psi(r/1)\Psi(1/d) = \psi(r)\psi(d)^{-1}$.
- But it is easy to see that this choice of Ψ does work, so it is the only such homomorphism.

- The point here is that $D^{-1}R$ is the smallest ring in which all elements of D become units.
- When D contains no zero divisors (which is automatically the case if R is a domain and D does not contain zero), then R injects into $D^{-1}R$ via $r \mapsto r/1$.
- A particular useful case of localization is to construct $\mathbb{Q}$ from $\mathbb{Z}$ (we take $D = \mathbb{Z}\backslash\{0\}$ and $R = \mathbb{Z}$) or more generally to construct the field of fractions of an integral domain R (take $D = R\backslash\{0\}$).

- We also note in passing that we can localize any R-module M in the same way: one simply writes down the same construction using pairs (m,d) with $m \in M$ and $d \in D$ in place of pairs (r,d).
- Alternatively, one can obtain the localization of an R-module using tensor products: $D^{-1}M \cong M \otimes_R D^{-1}R$. (This tensor product just extends scalars from R to $D^{-1}R$, which is exactly what $D^{-1}M$ is.)
- Exercise: Show that localization commutes with sums, intersections, quotients, finite direct sums, and is exact.
- Exercise: Show that if I is an ideal of R, then $D = R\backslash I$ is multiplicatively closed if and only if I is prime.

- Our main situation of interest is that of localizing at a prime: this is the case where R is an integral domain and $D = R\backslash P$ is the complement of a prime ideal P of R.
- Exercise: Show that if P is a prime ideal and $D = R\backslash P$, then $D^{-1}R$ is a local ring with unique maximal ideal $D^{-1}P = \pi(P) = e_P$, the extension of the ideal P to $D^{-1}R$.
- The utility of localizing at a prime is that it isolates the ring's behavior at that prime.
- Example: The localization of $\mathbb{Z}$ at the prime ideal (p) is the ring $\mathbb{Z}_{(p)} = \{a/b \in \mathbb{Q} : p \nmid b\}$ of rational numbers whose denominator is not divisible by p. Its unique maximal ideal is $p\mathbb{Z}_{(p)}$, the set of multiples of p. The quotient ring $\mathbb{Z}_{(p)}/p\mathbb{Z}_{(p)}$ is isomorphic to $\mathbb{Z}/p\mathbb{Z}$.
- Note that $\mathbb{Z}_{(p)}$ is not the ring of p-adic integers $\mathbb{Z}_p$: the p-adic integers are obtained by taking a completion of the localization $\mathbb{Z}_{(p)}$ under the p-adic metric (which we will define later).
- Example: Let k be a field and take R to be the ring of k-valued functions on a set S. If we let M_a be the set of functions vanishing at a point $a \in S$, then M_a is a maximal ideal of R. The localization $R_{M_a} = \{f/g \in R : g(a) \neq 0\}$ is the ring of k-valued rational functions defined at a. The unique maximal ideal of M_a is the ideal of all k-valued rational functions vanshing at a.

- This second example illustrates the utility of localizing at a prime, because it allows us to study the local behavior of a rational function near the point a.
- For example, the elements of M_a are precisely those rational functions vanishing at a, while the elements of M_a^2 are the rational functions that vanish to order 2 at a (i.e., have a double root), and so forth.
- More generally, if we localize a domain at a principal prime ideal, by looking at powers of the maximal ideal, we can measure what power of a prime a given element is divisible by.

- We now formalize all of this using discrete valuations, which provide us a way to identify primes using only the field structure:

- Definition: Let F be a field. A discrete valuation on F is a surjective function $v : F^\times \to \mathbb{Z}$ such that $v(ab) = v(a) + v(b)$ for all $a, b \in F$ and $v(a+b) \geq \min(v(a), v(b))$ for all $a, b \in F^\times$ with $a + b \neq 0$. The set $R = \{r \in F^\times : v(r) \geq 0\} \cup \{0\}$ is called the valuation ring of v.

  - For convenience, if v is a discrete valuation we often also write $v(0) = \infty$, in which case we can ignore the various exceptions in the definition above (e.g., $v(a+b) \geq \min(v(a), v(b))$ now holds for all a, b).
  - In general, we say an integral domain R is a discrete valuation ring (DVR) if it is the valuation ring for some discrete valuation on its field of fractions.
  - Example: For a fixed prime p, the p-adic valuation on $\mathbb{Q}$, which has $v_p(p^n \frac{r}{s}) = n$ for $p \nmid r, s$, is a discrete valuation. (For example, $v_2(4) = 2$, $v_2(\frac{1}{3}) = 0$, and $v_2(\frac{3}{4}) = -2$: the valuation simply gives the power of p in a rational number.) The associated valuation ring is the set of rational numbers whose denominator is not divisible by p: this is $\mathbb{Z}_{(p)}$, the localization of $\mathbb{Z}$ at (p).
  - Example: For a fixed irreducible polynomial p, the p-adic valuation on $\mathbb{F}_q(t)$, which has $v_p(p^n \frac{r}{s}) = n$ for $p \nmid r, s$, is a discrete valuation. (For example, $v_t(t^3) = 3$, $v_t(\frac{t}{t+1}) = 1$, and $v_{t+1}(\frac{t}{t+1}) = -1$.) The associated valuation ring is the set of rational functions whose denominator is not divisible by p: this is $A_{(p)}$, the localization of $A = \mathbb{F}_q[t]$ at (p).
  - In the two examples above, the valuation rings are both obtained as localizations. We can in fact construct DVRs by localizing in more generality.
  - Exercise (Corollary 8 from Section 16.2 of Dummit/Foote): If R is a Noetherian integrally-closed domain and P is a minimal nonzero prime ideal of R, then R_P is a DVR. Deduce in particular that if R is a Dedekind domain and P is a nonzero prime ideal, then R_P is a DVR.

- Proposition (Properties of DVRs): Let R be a discrete valuation ring with field of fractions F and valuation v. Also $t \in R$ be any element with $v(t) = 1$ (such an element is called a uniformizer). Then the following hold:

  1. For any $r \in F^\times$, either r or $1/r$ is in R.
  2. An element $u \in R$ is a unit of R if and only if $v(u) = 0$. In particular, if $\zeta \in F$ is any root of unity, then $v(\zeta) = 0$.
  3. If $x \in R$ is nonzero and $v(x) = n$, then x can be written uniquely in the form $x = ut^n$ for some unit $u \in R$.
  4. Every nonzero ideal of R is of the form (t^n) for some $n \geq 0$.
  5. The ring R is a Euclidean domain (hence also a PID and a UFD) and also a local ring.
  6. The ring S is a DVR if and only if it is a PID and a local ring but not a field.

     - Proofs: Exercises.

- We will also remark that a discrete valuation v on a field F naturally makes F into a metric space using the non-Archimedean metric $d_v(a, b) = 2^{-v(a-b)}$. Explicitly:

  1. We clearly have $d_v(a, b) \geq 0$ with equality if and only if $a = b$.
  2. Since $v(-1) = 0$ by (2) in the proposition above, we have $v(a - b) = v(b - a)$ and thus $d_v(a, b) = 2^{-v(a-b)} = 2^{-v(b-a)} = d_v(b, a)$.
  3. From $v(x+y) \geq \min(v(x), v(y))$ we have $v(a-b) \geq \min(v(a-c), v(c-b))$, so negating yields $-v(a-b) \leq \max(-v(a-c), -v(c-b))$. Then $d_v(a, b) = 2^{-v(a-b)} \leq \max(2^{-v(a-c)}, 2^{-v(c-b)}) = \max(d_v(a, c), d_v(c, b))$.

  - We could also replace 2 by any real number greater than 1 in the definition of the metric without affecting anything.

- With this metric, we can then speak fruitfully of Cauchy sequences, write down the metric topology on F, and take completions. (Completing $\mathbb{Q}$ under the p-adic metric yields the p-adic field $\mathbb{Q}_p$, while the completion of its valuation ring $\mathbb{Z}$ yields the ring of p-adic integers $\mathbb{Z}_p$.)

- We now have enough background to discuss primes in function fields. The point of all of these preliminaries is that there is a natural interplay between discrete valuations on F and the primes associated to F, at least in the case of $F = \mathbb{Q}$.

- Definition: If K is a function field over F, a prime P of K is the maximal ideal of a discrete valuation ring R containing F whose field of fractions is K. The associated valuation on K is denoted ord_P.

  - Explicitly, the idea is that if we have a discrete valuation on K, then the valuation ring R is a local ring whose unique maximal ideal represents a prime of K.

- It is worth going through why this definition is (up to some haziness) really the same as the usual one in the case $K = \mathbb{Q}$ that we already understand.

  - If we have a discrete valuation v on $\mathbb{Q}$, then $v(-1) = v(1) = 0$ and so $v(n) \geq 0$ for all integers n. This means that the valuation ring R contains $\mathbb{Z}$.
  - Exercise: If v is a discrete valuation on $\mathbb{Q}$, the set $P = \{n \in \mathbb{Z} : v(n) > 0\}$ is a prime ideal of $\mathbb{Z}$.
  - By the exercise, $P = (p)$ for some prime p. Then $v(a) = 0$ for $p \nmid a$, so if $v(p) = r$ we see $v(p^n \frac{a}{b}) = rn$ for $p \nmid a, b$. Since discrete valuations are onto and $v(p) > 0$, we must have $n = 1$, and so v is the usual p-adic valuation on $\mathbb{Q}$.
  - Therefore, the only discrete valuations on $\mathbb{Q}$ are the p-adic valuations. The corresponding valuation ring is then $\mathbb{Z}_{(p)}$ with unique maximal ideal $p\mathbb{Z}_{(p)}$.
  - We see that for each integer prime p, we obtain a unique prime ideal $P = p\mathbb{Z}_{(p)}$ inside the associated a valuation ring of $\mathbb{Q}$. The collection of valuations v_p, evaluated on a rational number $\alpha \in \mathbb{Q}$, measures "how divisible" the element α is by each of the primes p.
  - Note also that in this case, the quotient of the valuation ring $R = \mathbb{Z}_{(p)}$ by its maximal ideal $P = p\mathbb{Z}_{(p)}$ is isomorphic to $\mathbb{Z}/p\mathbb{Z}$, which has cardinality p: the size of this quotient R/P naturally gives us a way to measure the size of the prime P.

- We can do something quite similar in the function field case:

- Proposition (Degrees of Primes): If K is a function field over F and P is a prime with valuation ring R, then the quotient R/P is a finite-dimensional F-vector space. We define the degree of P to be the dimension of this vector space.

  - Proof: Since P is a maximal ideal of R and contains F, R/P is a field extension of F, so we just need to show its degree over F is finite.
  - Suppose $y \in P \backslash F$. As noted earlier, y is transcendental over F and $K/F(y)$ is a finite-degree extension. We claim that $[R/P : F] \leq [K : F(y)]$.
  - To see this, suppose that $x_1, x_2, \ldots, x_m \in R$ have the property that their reductions $\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m} \in R/P$ are F-linearly independent, and suppose there is a linear dependence over $F(y)$: say $f_1(y)x_1 + f_2(y)x_2 + \cdots + f_m(y)x_m = 0$ for some $f_i(y) \in F[y]$.
  - If we cancel any common factors of y from the $f_i(y)$, and then reduce modulo P, we obtain a linear dependence of the $\overline{x_i}$ in R/P, contradiction (the point is that not all of the f_i are divisible by y, so at least one of them has a nonzero reduction modulo P).
  - Thus, any linearly independent set in R/P lifts to a linearly independent set in K, so we obtain the claimed inequality.

- Let's work all of this out in the case we mostly understand already: the purely transcendental extension $K = F(t)$.

  - From our discussion, if p is any irreducible polynomial in $A = F[t]$, we obtain a discrete valuation ring associated to the prime p as the localization $R = A_{(p)}$ and its unique maximal ideal is $P = pA_{(p)}$.

- Then $R/P \cong A/(p)$, in which case the dimension of R/P as an F-vector space is the same as the dimension of $A/(p)$ as an F-vector space, and this is simply $\deg(p)$, since $\{\overline{1}, \overline{t}, \ldots, \overline{t^{\deg p - 1}}\}$ is a basis for $A/(p)$.
- Thus, the degree as defined above agrees perfectly with our normal sense of the degree of a polynomial.
- The associated p-adic valuation v_p is the same as the one we discussed earlier: $v_p(p^n \frac{r}{s}) = n$ for $p \nmid r, s$: it pick out the power of the prime p that divides a rational function $f = p^n \frac{r}{s} \in F(t)$.
- Localizing A at a prime ideal yields almost all of the possible discrete valuation rings attached to $F[t]$.
- But there is, in fact, one more: the valuation $v_\infty(f/g) = \deg(g)/\deg(f)$, whose associated valuation ring is obtained by localizing $A' = F[t^{-1}]$ at the prime ideal (t^{-1}). The resulting prime is known as the prime at infinity, and its degree is 1.
- Exercise: Prove that the p-adic valuations v_p along with v_∞ are the only discrete valuations on $F(t)/F$. (Use a similar argument to the one for $\mathbb{Q}$ by identifying all possible uniformizers.)
- The general philosophy is that there will be a few "infinite primes", and the rest are "finite primes" that arise from localizing at a prime ideal. (Over $\mathbb{Q}$, the infinite prime corresponds to the usual absolute value $|\cdot|$ resulting in the completion $\mathbb{R}$, but this is not a discrete valuation.)

- We can also give a brief explanation of some of the terminology (e.g., "function field").
  - Suppose P is a prime of K/F where F is algebraically closed (e.g., $F = \mathbb{C}$). Then the quotient R/P is a finite-degree field extension of F hence is simply (isomorphic to) F itself.
  - For any element $a \in K$, we can then simply "read off" the values of a at the various primes P by interpreting $a(P)$ as the image of a inside the quotient $R/P \cong F$.
  - This is why K/F is called a function field, since we may think of the actual elements of K as F-valued functions on the primes P. The elements of F correspond to constant functions, which is why we refer to F as the constant field of K.
  - Furthermore, as we will discuss later, we can think of the primes of P geometrically as "places" or "points".
  - For $K = \mathbb{C}(t)$, for example, we obtain a finite prime P_r corresponding to each element $t - r$ for $r \in \mathbb{C}$, along with the infinite prime. Explicitly, P_r is the collection of rational functions vanishing at r (which is the unique maximal ideal of the ring R of all rational functions defined at r), and the evaluation-at-r map yields an explicit isomorphism $R/P_r \cong \mathbb{C}$.
  - Together, the finite primes P_r along with the infinite prime form the complex projective line $\mathbb{P}^1(\mathbb{C}) = \mathbb{C} \cup \{\infty\}$, which we may view analytically as being the Riemann sphere, and the field K consists of all of the $\mathbb{C}$-valued rational functions on the Riemann sphere.

0.11 (Oct 8) Student Presentations of HW1

0.12 (Oct 15) Divisors and the Divisor Group

- Our main goal now is to state, show, and use the Riemann-Roch theorem, which is the most fundamental basic theorem about function fields. The first ingredient is divisors:
- Definition: The divisor group of K, written D_K, is the additive free abelian group generated by the primes of K.
  - The elements of D_K are of the form $D = \sum_P n_P P$ for $n_P \in \mathbb{Z}$, where all but finitely many of the n_P are zero. We will write $\mathrm{ord}_P(D) = n_P$.
- Definition: The degree of a divisor $D = \sum_P n_P P$ is $\deg(D) = \sum_P n_P \deg(P)$. The degree map is a homomorphism from D_K to $\mathbb{Z}$; its kernel is the set of degree-0 divisors.
  - Note that this sum is well defined since all but finitely many n_P are zero.
- If $a \in K^\times$ is nonzero, we can attach a divisor to it by calculating its order at each prime P of K.

- Definition: We define the divisor of an element $a \in K^\times$ as $\mathrm{div}(a) = \sum_P v_P(a)P$. The divisors of the form $\mathrm{div}(a)$ for some $a \in K^\times$ are called principal divisors.
    - We will often also write $\mathrm{ord}_P(a)$ (the order of a at P) interchangeably with $v_P(a)$ (the P-adic valuation of a).
    - Remark: In many sources, the divisor of a is often written (a). In our context, this can lead to ambiguities, since the same notation is also used for the ideal generated by a.
    - A priori, it is not clear we have actually given a well-defined divisor: to show this we need to establish that $v_P(a) = 0$ for all but finitely many primes P.
    - Assuming this for the moment, since $\mathrm{ord}_P(a/b) = \mathrm{ord}_P(a) - \mathrm{ord}_P(b)$, summing over all primes shows that $\mathrm{div}(a/b) = \mathrm{div}(a) - \mathrm{div}(b)$, so the principal divisors are a subgroup of the divisor group D_K.
- We must still show that the divisor of an element is actually well defined:
- Proposition (Divisors of Elements): For any $a \in K^\times$, we have $v_P(a) = 0$ for all but finitely many primes P of K.
    - Proof: First, if $a \in F^\times$, then for any prime P the associated valuation ring R contains F. In particular, since $a \in F^\times$ this means a is a unit in R hence has valuation 0. This means $v_P(a) = 0$ for all P, and so $\mathrm{div}(a) = 0$.
    - Now suppose $a \notin F^\times$, so a is transcendental over F and $K/F(a)$ is finite.
    - If P is a prime of K and $v_P(a) > 0$, then by definition there is a discrete valuation ring R such that a is not a unit. Then R contains $F[a]$, and since R is integral over $F[a]$, it embeds into the integral closure of $F[a]$ inside K.
    - We lose nothing by enlarging R, so now assume R is the integral closure of $F[a]$ in K: then R is a Dedekind domain[2] since it is Noetherian, integrally closed, and the localization of R at any nonzero prime is a discrete valuation ring.
    - Since R is a Dedekind domain, every nonzero ideal can be factored uniquely as a product of prime ideals. So write $Ra = \mathfrak{p}_1^{b_1}\mathfrak{p}_2^{b_2}\cdots\mathfrak{p}_k^{b_k}$ for distinct prime ideals $\mathfrak{p}_1, \mathfrak{p}_2, \ldots, \mathfrak{p}_k$. Localizing at the prime $\mathfrak{p}_i$ yields a unique prime P_i of K, and since all of the other ideals become invertible, we see $\mathrm{ord}_{P_i}(a) = b_i$, and $\mathrm{ord}_Q(a) \leq 0$ for any other prime Q.
    - In particular, we see that there are only finitely many primes for which $v_P(a) > 0$.
    - In the same way, if $v_P(a) < 0$, by doing the same calculation for a^{-1} (i.e., by taking R' to be the integral closure of $F[a^{-1}]$ inside K and factoring $R'a^{-1} = \mathfrak{q}_1^{c_1}\cdots\mathfrak{q}_l^{c_l}$ as a product of prime ideals) we see that there are also only finitely many primes for which $v_Q(a) < 0$.
    - Thus, the divisor $\mathrm{div}(a) = \sum_P v_P(a)P$ is well defined, as claimed.
- Note that the proof given above gives an algorithm for computing divisors of elements: inside the integral closure of $F[a]$ (and $F[a^{-1}]$) inside K, we compute the prime ideal factorization of the ideal (a) (or (a^{-1})); then the exponents of the various primes give the corresponding valuations.
    - Let's work this out in the case of $K = \mathbb{C}(t)$: suppose we have a nonconstant rational function $a = \dfrac{f}{g}$, which we factor as $a = u\dfrac{(t-r_1)^{a_1}\cdots(t-r_k)^{a_k}}{(t-s_1)^{b_1}\cdots(t-s_l)^{b_l}}$ for distinct $r_1, \ldots, r_k, s_1, \ldots, s_l \in \mathbb{C}$ and some unit factor $u \in \mathbb{C}^\times$.
    - Each of the monic irreducibles $t - r_i$ and $t - s_j$ yields a unique prime of K. We have $v_{P_{t-r_i}}(a) = a_i$ and $v_{P_{t-s_j}}(a) = -b_j$ for each i, j, and the valuation at every other finite prime is 0. Also, we have $v_\infty(a) = \deg(g) - \deg(f) = \sum b_j - \sum_i a_i$.
    - Therefore, $\mathrm{div}(a) = a_1P_{t-r_1} + \cdots + a_kP_{t-r_k} - b_1P_{t-s_1} - \cdots - b_lP_{t-s_l} + [\sum_j b_j - \sum_i a_i]\infty$.
    - Notice in particular that $\deg(\mathrm{div}(a)) = \sum_i a_i - \sum b_j + [\sum_j b_j - \sum_i a_i] = 0$, which is to say, the divisor of any element of $K^\times$ has degree zero.

[2] For additional reference about Dedekind domains, including the factorization result we quote here, see section 16.3 of Dummit/Foote.

- Furthermore, we see that the primes with positive order at a correspond precisely to the zeroes of $a = f/g$ (and the order of a at that prime is the order of vanishing of a there), while the primes with negative order correspond to poles (and the order of a at that prime is the order of the pole there).
- We can also give a similar calculation for $K = F(t)$ where F is not algebraically closed in terms of the monic irreducible factors of $a = f/g$. The resulting divisor decomposition is essentially just the prime factorization of the rational function f/g:
- Exercise: For $K = F(t)$, if $a = u\frac{p_1^{a_1}\cdots p_k^{a_k}}{q_1^{b_1}\cdots q_l^{b_l}}$ for $u \in F^\times$ and distinct monic irreducibles $p_1, \ldots, p_k, q_1, \ldots, q_l$ having associated primes $P_1, \ldots, P_k, Q_1, \ldots, Q_l$, show that $\operatorname{div}(a) = a_1P_1 + \cdots + a_kP_k - b_1Q_1 - \cdots - b_lQ_l + [\sum_j b_j \deg(q_j) - \sum_i a_i \deg(p_i)]\infty$.

- Motivated by the calculations for $K = \mathbb{C}(t)$, we can also pick out the zeroes (respectively, poles) of an element by extracting only the portion of its divisor with positive (respectively, negative) coefficients:
- Definition: If $a \in K^\times$ has divisor $\operatorname{div}(a) = \sum_P n_P P$, we define $\operatorname{div}_+(a) = \sum_P \max(0, n_P)P = \sum_{P:n_P>0} n_P P$ and $\operatorname{div}_-(a) = \sum_P \min(0, n_P)P = \sum_{P:n_P<0} n_P P$.
  - Notice that $\operatorname{div}(a) = \operatorname{div}_+(a) - \operatorname{div}_-(a)$ for any element a.
  - Remark: There are various other notations for these quantities that are often used, such as $(a)_0$ for div_+ and $(a)_\infty$ for div_-, which are intended to evoke the idea of picking out the zeroes and poles of a.
  - For $a = \frac{f}{g} = u\frac{(t-r_1)^{a_1}\cdots(t-r_k)^{a_k}}{(t-s_1)^{b_1}\cdots(t-s_l)^{b_l}}$ in $\mathbb{C}(t)$, we have $\operatorname{div}_+(a) = a_1P_1 + \cdots + a_kP_k$ (plus $[\deg g - \deg f]\infty$ if $\deg g - \deg f > 0$) and $\operatorname{div}_-(a) = b_1Q_1 + \cdots + b_lQ_l$ (plus $[\deg g - \deg f]\infty$ if $\deg g - \deg f < 0$).
  - Exercise: For any field F, if $f(t), g(t) \in F[t]$ are relatively prime, show that $[F(t) : F(\frac{f(t)}{g(t)})] = \max(\deg f, \deg g)$. [Hint: Use Gauss's lemma to show that $q(y) = f(y) - \frac{f(t)}{g(t)}g(y) \in F(\frac{f(t)}{g(t)})[y]$ is the minimal polynomial of t over $F(\frac{f(t)}{g(t)})$.]
  - In the example above, we can also compute that $\deg(\operatorname{div}_+(a)) = \deg(\operatorname{div}_-(a)) = \max(\deg f, \deg g)$, and by the exercise above, this quantity is equal to the extension degree $[K : F(a)]$. In fact, this result is true in general:

- Theorem (Divisor Degrees): For any $a \in K^\times$, we have $\deg(\operatorname{div}_+(a)) = \deg(\operatorname{div}_-(a)) = [K : F(a)]$. As a consequence, $\deg(\operatorname{div}(a)) = 0$.
  - We will defer the proof for $F = \mathbb{F}_q$ until later, since it requires a number of ingredients we have not developed yet. The general case we will skip (the result is not that difficult, but it is not especially enlightening for what we will be doing).
  - Our main observation here is that the divisor of an element $a \in K^\times$ always has degree 0, which is to say, the principal divisors are actually a subgroup of the group of degree-0 divisors.

- Definition: We say two divisors D_1 and D_2 are linearly equivalent (and write $D_1 \sim D_2$) if $D_1 - D_2$ is principal. The resulting equivalence classes (i.e., divisors modulo principal divisors) form a group called the class group, or the Picard group, of K.
  - Exercise: Verify that this relation is an equivalence relation and that the equivalence classes are the elements in the quotient group of divisors modulo principal divisors.
  - Some notation for these various groups: $\operatorname{Div}(K) = D_K$ is the group of all divisors on K, $\operatorname{Div}^0(K)$ is the group of degree-0 divisors on K, $\operatorname{Cl}(K) = \operatorname{Pic}(K) = \operatorname{Div}(K)/[\text{principal divisors}]$ is the class group of K.
  - Since principal divisors all have degree zero, we can also form the reduced Picard group $\operatorname{Pic}^0(K) = \operatorname{Div}(K)/[\text{principal divisors}]$.

- For $K = F(t)$, the reduced Picard group is trivial:
- Proposition (Reduced Picard Group of $F(t)$): If $K = F(t)$, then $\operatorname{Pic}^0(K) = \operatorname{Div}(K)/[\text{principal divisors}]$ is the trivial group, and $\operatorname{Pic}(K) \cong \mathbb{Z}$.

- Remark: It can be shown that $K = F(t)$ is essentially the only situation where $\mathrm{Pic}^0(K)$ is trivial.
- Proof: The result is equivalent to showing that every divisor of degree 0 is principal, so suppose $D = \sum_P b_P P$ has degree 0.
- Let $a = \prod_{P \neq \infty} p(t)^{b_P}$, where $p(t)$ is the irreducible polynomial associated to the finite prime P of K.
- Then $\mathrm{ord}_P(a) = b_p$ for each prime $P \neq \infty$. But since $\sum_P b_P \deg(P) = 0$ by the assumption on D, and $\deg(\mathrm{div}(a)) = 0$ as well, we must have $\mathrm{ord}_\infty(a) = b_\infty$ also.
- Then $\mathrm{ord}_P(a) = b_p$ for all primes P, meaning that $\mathrm{div}(a) = D$ and so D is principal as claimed.
- The statement that $\mathrm{Pic}(K) \cong \mathbb{Z}$ follows immediately from $\mathrm{Div}(K)/\mathrm{Div}^0(K) \cong \mathbb{Z}$.

- To finish the discussion here, we remark on the analogy with the case of algebraic number fields.
  - If $K/\mathbb{Q}$ is an algebraic number field, we have an exact sequence $1 \to [\text{units of } \mathcal{O}_K] \to K^* \to [\text{fractional ideals of } \mathcal{O}_K] \to [\text{ideal class group of } K] \to 1$.
  - If K/F is an algebraic function field, the analogous exact sequence is $1 \to F^* \to K^* \to \mathrm{Div}^0(K) \to \mathrm{Pic}^0(K) \to 1$.
  - The constant field of K plays the role of the units of an algebraic number field, the group of degree-0 divisors plays the role of the fractional ideals in the ring of integers, and the reduced Picard group plays the role of the ideal class group.

- We now put a partial ordering on divisors motivated by the idea of divisibility for integers and rational functions: the idea is that if we look at p-adic valuations of elements of $\mathbb{Q}$, we can identify the elements of $\mathbb{Z}$ as those whose valuations are nonnegative at every prime p.

- Definition: If a divisor $D = \sum_P n_P P$ has $n_P \geq 0$ for all primes P, we say D is effective and we write $D \geq 0$. We extend this notion to a partial ordering on divisors by writing $D_1 \leq D_2$ if and only if $D_2 - D_1$ is effective.
  - Exercise (easy): Check that the relation $D_1 \leq D_2$ is a partial ordering on divisors.
  - The partial ordering on divisors allows us to specify the order of zeroes and poles: to illustrate, for $K = \mathbb{C}(t)$, saying that f has a pole of order at most 2 at $z = 0$ and a zero of order at least 3 at $z = 1$ is equivalent to saying $\mathrm{div}(f) \geq -2P_{z-0} + 3P_{z-1}$.

- Definition: If D is a divisor, the Riemann-Roch space associated to D is the set $L(D) = \{a \in K^\times : \mathrm{div}(a) \geq -D\} \cup \{0\}$. Equivalently, an element $a \in K$ is in $L(D)$ if and only if $v_P(a) \geq -v_P(D)$ for all primes P of K.
  - When D is effective, $L(D)$ represents all rational functions whose poles are "no worse" than D.
  - More generally, if $D = \sum_P n_P P - \sum_Q m_Q Q$ with $n_i, m_i > 0$, then $L(D)$ consists of all $a \in K$ such that a has a zero of order at least m_Q at each prime Q, and may have poles only at the primes P, of order at most n_P at P.
  - It is not hard to see that $L(D)$ is an F-vector space: if $a, b \in L(D)$, then $a+b \in L(D)$ because $v_P(a+b) \geq \min(v_P(a), v_P(b))$ for each prime P, and $ca \in L(D)$ for all $c \in F$ since $v_P(ca) = v_P(c) + v_P(a) = v_P(a)$ since $v_P(c) = 0$ for all primes P.
  - Example: For $K = F(t)$ and $D = P_t$, we can see that $L(D) = \mathrm{span}(1, t^{-1})$, since the only possible poles of an element $f/g \in L(D)$ function occur at $t = 0$ (of order 1) and the function must also have $\deg g \geq \deg f$ since there is no pole at the infinite prime P_∞.
  - Example: For $K = F(t)$ and $D = 3P_\infty$, we can see that $L(D) = \mathrm{span}(1, t, t^2, t^3)$ since the function f/g has no poles except a pole of order at most 3 at P_∞ (meaning that $\deg g \leq \deg f + 1$), which is to say, f/g is a polynomial of degree at most 3.
  - Example: For $K = F(t)$ and $D = -P_t$, we can see that $L(D) = \{0\}$, since any nonzero element $f/g \in L(D)$ would need to be zero at $t = 0$ and defined at all other primes, but this cannot occur because g would have to be constant, but then $\deg f > \deg g$ would force f/g to have a pole at P_∞.
  - Example: For arbitrary K, we have $L(0) = F$, since $\mathrm{div}(a) = 0$ for all $a \in F^\times$, but any element $x \in K^\times \backslash F$ necessarily has at least one pole (at any prime associated to a prime occurring in the prime factorization of the ideal generated by x^{-1} in the integral closure of $F[x^{-1}]$ inside K).
  - Exercise: Determine $L(D)$ when $K = F(t)$ for $D = P_t - P_\infty$, $P_t + P_\infty$, and $P_t + P_{t-1}$.

0.13 (Oct 20) The Riemann-Roch Theorem and Applications

- We now study the dimensions of these Riemann-Roch spaces.
- Definition: If D is a divisor, we define $\ell(D) = \dim_F L(D)$.
  - Examples: By the examples worked out above, for $K = F(t)$ we have $l(P_t) = 2$, $l(3P_\infty) = 4$, and $l(-P_t) = 0$.
  - Example: For an arbitrary K, we have $\ell(0) = 1$, since $L(0) = F$.
- Proposition (Properties of $l(D)$): Let K be a function field over F and D be a divisor of K.
  1. If $D_1 \leq D_2$, then $\ell(D_1) \leq \ell(D_2)$.
     - Proof: This follows immediately from the definition, since $D_1 \leq D_2$ clearly implies that $L(D_1)$ is a subspace of $L(D_2)$.
  2. If $D_1 \sim D_2$, then $L(D_1) \cong L(D_2)$ and so $\ell(D_1) = \ell(D_2)$.
     - Proof: If $D_1 = D_2 + \mathrm{div}(g)$, then the map from $L(D_1)$ to $L(D_2)$ sending $f \mapsto fg$ is easily seen to be an isomorphism of vector spaces since it has an inverse map $h \mapsto h/g$.
  3. If $\deg(D) \leq 0$, then $L(D) = \{0\}$ and $l(D) = 0$ except when $D = \mathrm{div}(a)$ is principal, in which case $L(D) = \mathrm{span}(a)$ and $l(D) = 1$.
     - Proof: Suppose $f \in L(D)$ and $f \neq 0$. Then $0 = \deg(\mathrm{div}(f)) \geq \deg(-D) = -\deg(D)$.
     - Furthermore, equality can hold only if $D = -\mathrm{div}(f)$ for some $f \in K^\times$, in which case D is principal.
     - If D is principal, then $\ell(D) = \ell(0) = 1$ by (2), and $L(D) = Fa = \mathrm{span}(a)$ by the same calculation.
  4. If D_1 and D_2 are divisors with $D_1 \leq D_2$, then $\dim_F(L(D_2)/L(D_1)) \leq \deg(D_2) - \deg(D_1)$.
     - Proof: Induct on the sum of the coefficients of the primes in the effective divisor $D_2 - D_1$. The base case $D_2 - D_1 = 0$ is trivial.
     - For the inductive step, suppose that $D_2 = D_1 + P$ for some prime P, and choose $x \in K$ such that $v_P(x) = v_P(D_2) = v_P(D_1) + 1$.
     - Then for any $y \in L(D_2)$, we have $v_P(xy) = v_P(x) + v_P(y) \geq v_P(D_2) - v_P(D_2) \geq 0$, so $xy \in R_P$ where R is the valuation ring associated to the prime P.
     - By composing with the evaluation map at P (i.e., taking the quotient of R_P by PR_P and then viewing this as isomorphic to R/P), we obtain an F-linear transformation $\varphi : L(D_2) \to R_P/PR_P \cong R/P$ with $\varphi(y) = (xy)(P)$.
     - Then $y \in \ker(\varphi)$ if and only if $(xy)(P) = 0$ if and only if $v_P(xy) \geq 1$ if and only if $v_P(y) \geq 1 - v_P(D_2) = -v_P(D_1)$, and this last statement is equivalent to $y \in L(D_1)$.
     - Thus, by the first isomorphism theorem, we have an injection from $L(D_2)/L(D_1)$ to R/P. Taking dimensions yields $\dim_F(L(D_2)/L(D_1)) \leq \dim_F(R/P) = \deg(P)$.
     - This establishes the inductive step, so the general result follows.
  5. For any effective divisor D, we have $\ell(D) \leq \deg(D) + 1$. In fact, this inequality holds for any divisor D of degree ≥ 0.
     - Proof: For effective divisors, this follows immediately by induction on the degree of D using (4), starting with the base case $l(0) = 1$.
     - For general divisors, the result is trivial if $\ell(D) = 0$, so suppose otherwise that $\ell(D) \geq 1$ and let $a \in L(D)$ be nonzero. Then $\mathrm{div}(a) \geq -D$ which is equivalent to $D - \mathrm{div}(a^{-1}) \geq 0$.
     - Then for $D' = D - \mathrm{div}(a^{-1})$, we see that D is equivalent to the effective divisor D', and so by (2) we have $\ell(D) = \ell(D') \leq \deg(D') + 1 = \deg(D) + 1$, as required.
  6. For any divisor D, the quantity $\ell(D)$ is finite.
     - Proof: If $\deg(D) < 0$ then (3) gives $\ell(D) = 0$, while if $\deg(D) \geq 0$ then (5) gives $\ell(D) \leq \deg(D) + 1$.
- What we would like to be able to do now is to calculate the actual dimension $\ell(D)$ for arbitrary divisors D. Rather than delaying the point, we will now get right to our main result:

- Theorem (Riemann-Roch): For any algebraic function field K/F, there exists an integer $g \geq 0$, called the genus of K, and a divisor class $\mathcal{C}$, called the canonical class of K, such that for any divisor $C \in \mathcal{C}$ and any divisor $A \in \mathrm{Div}(K)$, we have $\ell(A) = \deg(A) - g + 1 + \ell(C - A)$.
    - Remarks: The divisor class $\mathcal{C}$, as we will explain at length later in the case of Riemann surfaces, is the divisor class associated with the Weil differentials of K.
- We will not prove the general function-field version of the Riemann-Roch theorem, since it requires a fair bit of background to develop the necessary results about Weil differentials.
    - Instead, we will go through the proof of the analytic version of Riemann-Roch for Riemann surfaces, which contains most of the main ideas but is more accessible since the complex-analytic notion of a differential is quite natural.
- For now, we will run through some consequences of the Riemann-Roch theorem.
- Proposition (Corollaries of Riemann-Roch): Let K/F be an algebraic function field.
    1. For any divisor A with $\deg(A) \geq 0$, we have $\deg(A) - g + 1 \leq \ell(A) \leq \deg(A) + 1$.
        - Proof: We showed the upper bound earlier using an inductive argument. The lower bound follows immediately from Riemann-Roch since $\ell(C - A) \geq 0$.
    2. For $C \in \mathcal{C}$ we have $\ell(C) = g$ and $\deg(C) = 2g - 2$.
        - Proof: First set $A = 0$ in Riemann-Roch: this yields $\ell(0) = \deg(0) - g + 1 + \ell(C)$, so since $\ell(0) = 1$ and $\deg(0) = 0$, we get $\ell(C) = g$.
        - Now set $A = C$ in Riemann-Roch: this yields $\ell(C) = \deg(C) - g + 1 + \ell(0)$, and so $\deg(C) = \ell(C) + g - 1 - \ell(0) = 2g - 2$.
    3. If $\deg(A) \geq 2g - 2$, then $\ell(A) = \deg(A) - g + 1$ except when $A \in \mathcal{C}$ (in which case $\ell(A) = g$).
        - Proof: If $\deg(A) \geq 2g - 2$, then $\deg(C - A) \leq 0$, and so $\ell(C - A) = 0$ except when $C - A$ is principal (i.e., when $A \in \mathcal{C}$).
        - When $\ell(C - A) = 0$ Riemann-Roch immediately gives $\ell(A) = \deg(A) - g + 1$, and when $A \in \mathcal{C}$ we have $\ell(A) = g$ by (2).
    4. The genus g is unique, as is the equivalence class $\mathcal{C}$.
        - Proof: Pick A of sufficiently large degree: then $\deg(A) - \ell(A) + 1 = g$ by (3), so g is uniquely determined.
        - For $\mathcal{C}$, if $\ell(A) = \deg(A) - g + 1 + \ell(C - A) = \deg(A) - g + 1 + \ell(D - A)$ for some other divisor D, then $\ell(C - A) = \ell(D - A)$ for all A.
        - Setting $A = C$ yields $\ell(D - C) = 1$ and setting $A = D$ yields $\ell(C - D) = 1$, and these are contradictory unless $D - C$ is principal, which is to say, $D \sim C$.
- Let's use Riemann-Roch to examine function fields of small genus. We start with the simplest genus $g = 0$.
    - By Riemann-Roch, we have $\ell(A) = \deg(A) + 1 + \ell(C - A)$ for any divisor A, and also $\deg(C) = -2$.
    - Also, by (3), if $\deg(A) \geq -1$ then $\ell(A) = \deg(A) + 1$. In particular, since $\deg(-C) = 2$, we have $\ell(-C) = 3$.
    - Now, for any prime P, we have $\ell(P) \leq \deg(P) + 1$. So, if P is any prime with $P \leq C$ (there must be at least one since $\deg(-C)$ is positive), we see $\ell(P) \leq \ell(-C) = 3$. Thus, $\deg(P)$ must be either 1 or 2.
    - First suppose that there is a prime P of degree 1. Then $\ell(P) = 2$. Since F is a subspace of $L(P)$, there is a basis of $L(P)$ of the form $\{1, x\}$ for some $x \notin F$.
    - Then since $\deg(\mathrm{div}(x) + P) = 1$ and $\mathrm{div}(x) + P \geq 0$, we must have $\mathrm{div}(x) + P = Q$ for some prime Q (necessarily of degree 1). Then $\mathrm{div}(x) = P - Q$, and so $[K : F(x)] = \deg(\mathrm{div}_+(x)) = \deg(P) = 1$, which means $K = F(x)$.
    - Exercise: Show in this case that the canonical class contains every divisor of K of degree -2.
    - Now suppose there is no prime P of degree 1: per earlier, we have a prime $P \leq C$ of degree 2.

- Then $\ell(P) = 3$, so again since $L(P)$ contains F, we may take a basis for $L(P)$ of the form $\{1, x, y\}$ for some F-linearly independent $x, y \notin F$.
- In the same way as above, we see that $\text{div}(x) = P - Q$ and $\text{div}(y) = P - R$ for some (necessarily distinct) primes Q and R of degree 2.
- Then $[K : F(x)] = \deg(\text{div}_+(x)) = 2$ and $[K : F(y)] = \deg(\text{div}_+(y)) = 2$ also. Since $F(x) \neq F(y)$ (by linear independence and the fact that K is a degree-2 extension of both), we see $K = F(x, y)$.
- Furthermore, Riemann-Roch says that $\ell(2P) = 1 + \deg(2P) = 5$, but we can find six different elements in $L(2P)$, namely $\{1, x, y, x^2, xy, y^2\}$. They must therefore be F-linearly dependent, so we see that x and y satisfy some quadratic relation $ax^2 + bxy + cy^2 + dx + ey = f$.
- Geometrically, this case corresponds to a conic, while the case $K = F(x)$ corresponds to a line (since we can think of $F(x) = F(x, y)$ where y is a linear function of x).

- We can use similar ideas to study the case where the genus g is equal to 1.
  - In this case, for $g = 1$ Riemann-Roch and its corollaries say that $\ell(A) = \deg(A) + \ell(C - A)$, that $\deg(C) = 0$ and $\ell(C) = 1$, and that if $\deg(A) \geq 1$ then $\ell(A) = \deg(A)$.
  - Unlike the case $g = 0$, we are not necessarily guaranteed to have a prime of any given degree any more, since we cannot use C to construct a prime of small degree – indeed, since $\deg(C) = 0$ and $\ell(C) = 1$, in fact C is principal (and $C \sim 0$).
  - So let us instead merely suppose that we do have a prime P of degree 1. Then $\ell(2P) = 2$, so choose a basis $\{1, x\}$ for $L(2P)$, where we necessarily must have $v_P(x) = 2$ since $x \notin L(P)$. Then $\ell(3P) = 3$, so choose a basis $\{1, x, y\}$ for $L(3P)$, where we must necessarily have $v_P(y) = 3$ since $y \notin L(2P)$.
  - Then, as above, $[K : F(x)] = \deg(\text{div}_+(x)) = 2$ and $[K : F(y)] = \deg(\text{div}_+(y)) = 3$, so since 2 and 3 are relatively prime, we see $K = F(x, y)$.
  - Now we would like to identify what kind of algebraic relation x, y must satisfy (they are, after all, algebraically dependent), which we can do by looking at the spaces $L(kP)$ for larger values of k, since the various monomials $x^i y^j$ will all only have poles at P.
  - We have $\ell(4P) = 4$, but we can only identify 4 elements that must lie in this space: $\{1, x, y, x^2\}$. (In fact, they are all linearly independent since they all have different valuations at P.)
  - Likewise, $\ell(5P) = 5$, but we only have 5 elements in this space: $\{1, x, y, x^2, xy\}$, but again, these elements are all linearly independent since they have different valuations at P.
  - But now consider $\ell(6P) = 6$: we can generate 7 elements in this space: $\{1, x, y, x^2, xy, x^3, y^2\}$. We must therefore have a linear dependence among these elements, and in fact since x^3 and y^2 are the only elements with valuation 6 at P, they must both occur with nonzero coefficients.
  - By rescaling x, y appropriately, we obtain an algebraic relation of the form $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$ for some $a_1, a_2, a_3, a_4, a_6 \in F$: this is an elliptic curve, and the corresponding function field $K = F(x, y)$ is called an elliptic function field. (The indices on the coefficients a_i are listed that way because they are giving the "missing" pole valuation at P for the corresponding monomial term.)
  - When the characteristic of F is not 2 or 3, we may complete the square in y and the cube in x to obtain a simpler equation $y^2 = x^3 + Ax + B$.

- Our analysis, even in these comparatively simple situations, indicates a correspondence between algebraic function fields and algebraic plane curves.
  - As we will discuss later at length, there is an equivalence of the following two categories:
    1. (Objects) Algebraic function fields K/F of transcendence degree 1 where $K \cap \overline{F} = F$
       (Morphisms) Field injections fixing 1 (up to isomorphism)
    2. (Objects) Smooth projective curves defined over F
       (Morphisms) Non-constant rational maps defined over F (up to isomorphism)
  - The correspondence is obtained by associating a smooth curve C with the field of rational functions defined on C.

- We will now use Riemann-Roch in our specific case of interest with base field $F = \mathbb{F}_q$. Our first result is that there necessarily exist divisors of all degrees:
- Proposition (Existence of Degree-1 Divisors): If K is a function field over $\mathbb{F}_q$, then there exists a divisor D of degree 1 over K, and hence there exist degrees of all degrees over K.
  - Proof (sketch): Let P be a prime of K and let $\sigma = \text{Frob}_q$ be the q-power Frobenius automorphism of K.
  - Exercise: If R is the valuation ring of P, show that σR is also a valuation ring with maximal ideal σP, and that σ gives an isomorphism of R/P with $\sigma R/\sigma P$. Show also that for any $a \in K$, $v_{\sigma P}(a) = v_P(\sigma^{-1}a)$.
  - By the exercise, σP is also a prime of K and it has the same degree as P, so $\sigma P - P$ has degree 0.
  - It can be shown that $\sigma P - P$ is equal to $\sigma D - D$ for some degree-0 divisor D (this is essentially Hilbert's theorem 90), which means $\sigma(P - D) = P - D$.
  - Then $P - D$ is a divisor that is fixed by the Frobenius map, which (one may show) necessarily implies that $P - D$ has degree 1. (The principle is the same as the observation that the elements of $\overline{\mathbb{F}_q}$ fixed by Frobenius are precisely the elements of $\mathbb{F}_q$, which generate extensions of degree 1.)
- As an immediate consequence of the proposition above, we obtain an exact sequence $0 \to \text{Pic}^0(K) \to \text{Pic}(K) \to \mathbb{Z} \to 0$.
  - Our next goal is to show that the reduced Picard group is finite.
- Proposition (Finiteness of the Class Group): Let K be a function field over $F = \mathbb{F}_q$. Then the following hold:

1. For any $n \geq 1$, the number of primes of K having degree n is finite.
   - Proof: Let $x \in K \backslash F$, so that $[K : F(x)]$ is finite. If P is a prime of K having degree n, then P lies over some prime of $F(x)$ of degree $\leq n$.
   - Since there are only finitely many primes in $F(x)$ of degree $\leq n$, and there are only finitely many different primes in a finite-degree extension $[K : F(x)]$ that lie above a particular prime in $F(x)$ (specifically, this number is bounded by the extension degree), we see that there are only finitely many primes of K.
   - Exercise: Show that the number of primes of degree $\leq n$ in K is at most $[K : F(x)]q^n$ for any $x \in K \backslash F$.
2. For any $n \geq 0$, the number of effective divisors of K having degree n is finite.
   - Proof: Suppose $D = \sum_P n_P P$ is effective and has degree n. Then $\deg(P) \leq \deg(D) = n$ for each prime P appearing with a positive coefficient.
   - By (1), there are only finitely many possible primes P of degree at most n. For each such prime P, the coefficient n_P is at most $n/\deg(P)$, so there are finitely many possible choices for each n_P.
   - Thus, there are only finitely many possible terms $n_P P$ that can appear in D, and so the number of effective divisors of degree n is finite.
   - Exercise: Give an explicit upper bound in terms of $[K : F(x)]$, q, and n for the number of effective divisors of degree n in K.
3. The reduced Picard group $\text{Pic}^0(K)$ is finite.
   - Proof: Let D be a divisor of degree 1. If A is any divisor of degree 0, then $\deg(gD + A) = g$, so $\ell(gD + A) \geq \deg(gD + A) - g + 1 = 1$ by Riemann-Roch.
   - Pick any nonzero $f \in L(gD + A)$: then $\text{div}(f) + gD + A \geq 0$ is some effective divisor B.
   - Then $A \sim B - gD$, so since there are finitely many possible B by (2), and gD is fixed, there are finitely many possible classes for A.
4. If h_K is the class number of K (the cardinality of $\text{Pic}^0(K)$), then there are exactly h_K divisor classes of each possible degree.
   - Proof: Let D be a divisor of degree 1. Then for any divisor A of degree n, we see $A - nD$ has degree 0, and so by (3) there are h_K possible classes for $A - nD$ up to equivalence.
   - Since nD is fixed, this means there are h_K possible classes for A up to equivalence, as required.

0.14 (Oct 22) Proof of Riemann-Roch Over $\mathbb{C}$

- We now discuss the details of Riemann-Roch in the case $F = \mathbb{C}$, where we can give an essentially complete argument (aside from some reliance on a few facts from complex analysis and differential topology).
  - Under the correspondence of curves and function fields, we are analyzing smooth complex projective curves, which are the same as 1-dimensional complex differentiable manifolds.
  - If we instead work over the reals, we can equivalently think of a 1-dimensional complex differentiable manifold as a compact Riemann surface X.
  - In this situation, the genus g also represents the topological genus of X (i.e., the number of "holes" in the surface, also equal to $1 - \chi/2$ where χ is the Euler characteristic).
  - The primes of the function field K are then simply the points P in X (since we are over $\mathbb{C}$, all of the primes have degree 1).
  - The elements of the function field K are then the meromorphic functions on X (i.e., the functions that are complex-differentiable except at a finite set of poles).
  - For $f \in K^\times$, the divisor $\operatorname{div}(f) = \sum_{P \in X} v_P(f)P$ tabulates the zeroes and poles of f: if $v_P(f) = k > 0$ then f has a zero of order k at P, while if $v_P(f) = -k < 0$ then f has a pole of order k at P.
  - Two divisors D_1 and D_2 are equivalent when $D_1 - D_2$ is principal, which is to say, when they differ by the divisor of a meromorphic function.
  - We can also deduce a few facts about divisors of functions analytically (rather than algebraically as we did earlier).
  - For example, suppose $\operatorname{div}(f) = 0$: then f has no poles and is therefore holomorphic, but since X is compact this means $|f|$ is bounded and so by Liouville's theorem, f is constant. (This also shows that the only functions holomorphic on all of X are the constants.)
  - Exercise: For any nonzero meromorphic f on X, show that $\deg(\operatorname{div}(f)) = 0$. [Hint: Use Cauchy's argument principle: for any contour C, $\frac{1}{2\pi i}\int_C \frac{f'}{f}\,dz = Z - P$ is the number of zeroes minus the number of poles in C.]
- So far all of the basic theory is the same. However, on a differentiable manifold, we also have a natural notion of a meromorphic differential $d\omega$.
  - Specficially, a meromorphic 1-form (also called a meromorphic differential) is a differential that may locally be written as $d\omega = f(z)\,dz$ for some meromorphic function f, where z is the local coordinate. (Being more precise requires a careful discussion of local coordinates and charts.)
  - Example: If X is the Riemann sphere with its usual coordinate z on $\mathbb{C}$ and $1/z$ on $\mathbb{C} \cup \{\infty\}\backslash\{0\}$, some examples of meromorphic differentials are $z\,dz$ (it has a zero at 0 and a pole at ∞) and $\frac{z}{z+1}\,dz$ (it has a zero at 0 and a pole at -1).
  - We can then define the divisor of a meromorphic differential $d\omega = f\,dz$ as $\operatorname{div}(d\omega) = \sum_{P \in X} v_P(f)P$.
  - If $d\omega_1$ and $d\omega_2$ are two meromorphic differentials, then $d\omega_1/d\omega_2 = f_1/f_2$ is locally a ratio of meromorphic functions hence is itself a meromorphic function.
  - This means all meromorphic differentials share the same divisor class: this is the canonical class $\mathcal{C}$.
  - We also have the natural notion of a holomorphic differential, which is a meromorphic differential having no poles. (This is the differential analogue of an effective divisor.)
  - From differential topology, we have the following fundamental fact: the dimension of the vector space of holomorphic differentials on X is equal to the genus g. (Very roughly speaking, we can obtain independent holomorphic differentials by integrating around non-contractible paths on X.)
  - Exercise: Explain why this fact is equivalent to saying $\ell(C) = g$.
- Our goal now is to give a concrete way to understand the dimension $\ell(C - A)$ for a divisor A.
  - We can do this by defining a space of differentials that mimics the Riemann-Roch space $L(D) = \{a \in K : \operatorname{div}(a) \geq -D\}$.

- Definition: If D is a divisor on a compact Riemann surface X, we define $\Omega(D)$ to be the space of differentials $d\zeta$ such that $\operatorname{div}(d\zeta) \geq D$.
  - In the same way as for $L(D)$, it is easy to see that $\Omega(D)$ is an F-vector space.
  - Note that $\Omega(0)$ is the space of holomorphic differentials on X, which has dimension g as we noted earlier.
- Proposition: For any divisor D on a compact Riemann surface X, if $C = \operatorname{div}(\omega)$ for a meromorphic differential ω, then $\Omega(D) \cong L(C - D)$.
  - Proof: Suppose that $d\zeta \in \Omega(D)$ and consider $d\zeta/d\omega$: it is some meromorphic function f, and we have $\operatorname{div}(f) = \operatorname{div}(d\zeta/d\omega) = \operatorname{div}(d\zeta) - \operatorname{div}(d\omega) \geq D - C$, so $f \in L(C - D)$. (Remember that the definition for the Riemann-Roch space L has a minus sign!)
  - Thus, the map $d\zeta \mapsto d\zeta/d\omega$ is a linear transformation from $\Omega(D)$ to $L(C-D)$, and since clearly the map $f \mapsto f\,d\omega$ is an inverse, it is an isomorphism.
- We now have most of the necessary ingredients for Riemann-Roch. The key additional piece is to introduce the calculation of residues of functions and differentials at a point P.
  - Given a nonzero rational function f, we may write f as a Laurent series centered at P: i.e., as $f = \sum_{n=k}^{\infty} a_k z^k$ where $k = v_P(f)$ (which may be positive or negative) and z is the local uniformizer at P. We define the residue of f at P to be the coefficient a_{-1}.
  - By Cauchy's residue theorem, we can also calculate residues via integration: $\int_C f(z)\,dz = 2\pi i \sum_P \operatorname{Res}_P(f)$, where the sum is over all points P inside the contour C. In particular, by reversing the orientation of the curve and summing the results, we can see that the sum over all P of the residues of f is zero. (This is essentially just Stokes's theorem.)
  - In particular, if we have an effective divisor $D = P_1 + P_2 + \cdots + P_k$ for distinct points P_i, we obtain a map $\varphi : L(D) \to \mathbb{C}^k$ by taking $\varphi(f) = (\operatorname{Res}_{P_1}(f), \operatorname{Res}_{P_2}(f), \ldots, \operatorname{Res}_{P_k}(f))$. The kernel of this map is the set of functions $g \in L(D)$ whose residue is zero at each P_i, but this would mean g is holomorphic on all of X, hence constant.
  - Thus, we obtain an exact sequence $0 \to \mathbb{C} \to L(D) \xrightarrow{\varphi} \mathbb{C}^k$.
- Intuitively, the statement of Riemann-Roch now comes from trying to answer the question: how close is the map φ to being surjective? In other words, what conditions are there on the values of the residues of a meromorphic function in $L(D)$ at the points P_i?
  - We can answer this question by looking at the residues of holomorphic and meromorphic differentials.
  - If $d\omega \in \Omega(0)$ is holomorphic, we define the residue of $d\omega$ at P as the residue of the ratio $\dfrac{d\omega}{dz}$ at P where dz is the local uniformizer at P (this is well-defined because $\dfrac{d\omega}{dz}$ is a meromorphic function).
  - In the same way as for functions, the sum of the residues of any meromorphic differential over all points must be zero: thus, for each $d\omega \in \Omega(0)$ and $f \in L(D)$, we see that the sum of the residues of $f\,d\omega$ must be zero. This means each differential imposes a linear condition on the possible choices of residues for f.
  - More precisely, if $D = P_1 + P_2 + \cdots + P_k$ for distinct points P_i, we obtain a map $\psi : \Omega(0) \to \mathbb{C}^k$ by taking $\psi(D) = (\operatorname{Res}_{P_1}(d\omega), \operatorname{Res}_{P_2}(d\omega), \ldots, \operatorname{Res}_{P_k}(d\omega))$. The kernel of this map is the set of differentials $d\omega \in \Omega(0)$ whose residue is zero at each P_i, which is to say, $d\omega \in \Omega(D)$.
  - Thus, we obtain another exact sequence $0 \to \Omega(D) \to \Omega(0) \xrightarrow{\psi} \mathbb{C}^k$.
  - The images of the two maps φ and ψ are orthogonal by the observation made above: for any $f \in L(D)$ and any $\omega \in \Omega(0)$, the inner product of $\varphi(f)$ and $\psi(d\omega)$ is $\sum_{i=1}^k \operatorname{Res}_{P_i}(f) \cdot \operatorname{Res}_{P_i}(d\omega) = \sum_{i=1}^k \operatorname{Res}_{P_i}(f\,d\omega) = 0$ by Stokes's theorem.
  - So, since the images of φ and ψ are orthogonal, we see that $\dim(\operatorname{im}\varphi) + \dim(\operatorname{im}\psi) \leq k = \deg(D)$.
  - By the nullity-rank theorem, since $\ker(\varphi) = \mathbb{C}$ we get $\dim(\operatorname{im}\varphi) = \dim(L(D)) - 1 = \ell(D) - 1$.
  - Likewise, since $\ker(\psi) = \Omega(D)$ we get $\dim(\operatorname{im}\psi) = \dim(\Omega(0)) - \dim(\Omega(D)) = g - \ell(C - D)$.
  - Thus, we obtain the inequality $\ell(D) - 1 + g - \ell(C - D) \leq \deg(D)$.

- If we had equality everywhere (i.e., if the images of φ and ψ were actually orthogonal complements) then we would get the Riemann-Roch theorem!
    - As it is, we only have the weaker statement that $\ell(D) - 1 + g - \ell(C - D) \leq \deg(D)$, which is known as Riemann's inequality (and only in the case where D is effective and a sum of distinct points $P_1 + \cdots + P_k$, though we can remove the restriction that D is a sum of distinct points using a limiting argument).
    - One can in fact establish that the images of φ and ψ are orthogonal complements with quite a bit more work.
    - In the event that $C - D$ is also effective, however, we can extract the desired result just from Riemann's inequality: in such a case, we have $\ell(D) - 1 + g - \ell(C - D) \leq \deg(D)$ and also $\ell(C - D) - 1 + g - \ell(D) \leq \deg(C - D) = \deg(C) - \deg(D)$, so adding the two inequalities yields $2g - 2 \leq \deg(C)$.
    - But since $\deg(C) = 2g - 2$ (another calculation we take for granted), we must have equality in both cases.
    - This establishes Riemann-Roch for divisors D where both D and $C - D$ are effective divisors (or equivalent to effective divisors, since as we showed, $\ell(D_1) = \ell(D_2)$ when $D_1 \sim D_2$).
    - In fact, this is nearly enough to get the general result, since as we showed, if $L(D) \neq 0$ then D is equivalent to an effective divisor. In general, one needs to verify that when $\ell(C - D) = 0$, one has $\deg(D) \geq \ell(D) - 1 + g$.
    - Assuming the inequality $\deg(D) \geq \ell(D) - 1 + g$, one obtains the general statement of Riemann-Roch: if both D and $C - D$ are equivalent to effective divisors, the result is as above, and if D is but $C - D$ is not, the result follows from $\deg(D) \geq \ell(D) - 1 + g$, and if $C - D$ is but D is not, the result is equivalent by interchanging D and $C - D$.
    - Finally, if neither D nor $C - D$ is equivalent to an effective divisor (i.e., if $\ell(D) = \ell(C - D) = 0$), then by the inequality above we must have $\deg(D) \geq g - 1$ and $\deg(C - D) \geq g - 1$. But since $\deg(C) = 2g - 2$ this forces $\deg(D) = g - 1$, in which case we do get $\deg(D) = \ell(D) - 1 + g - \ell(C - D)$, as required.
- To summarize, the main tools used in proving Riemann-Roch involve studying the relationships between divisors and differentials, and using structural statements about residues of functions and differentials. In order to prove Riemann-Roch in an arbitrary function field K, we would need to develop the analogues of all of these ingredients over K.
    - We already have divisors and can define the residue of an element of K at a prime P by using series expansions in terms of uniformizers, similarly to how it works generally in $\mathbb{C}$.
    - Explicitly, given a nonzero element $a \in K^\times$ and a prime P, we can express a as a Laurent series with coefficients in the residue field R/P with respect to a fixed uniformizer t at P.
    - For simplicity first suppose that $v_P(a) \geq 0$: then $a \in R$ is defined at P. Let $\varphi : R \to R/P$ be the projection map, by definition the "value" of a at P is $\varphi(a) = a_0$; then $a - a_0 \in P$ so $a = a_0 + b_1 t$ for an element $b_1 \in R$.
    - Now let $\varphi(b_1) = a_1$, so $b_1 - a_1 \in P$ and thus $b_1 = a_1 + b_2 t$ for some $b_2 \in R$: now $a = a_0 + a_1 t + b_2 t^2$. We may clearly continue this process indefinitely to generate a power series expansion $a = a_0 + a_1 t + a_2 t^2 + \cdots$.
    - Exercise: Let P be a prime of the function field K with valuation ring R and residue field $E = R/P$. Show that the power series expansion method yields a one-to-one ring homomorphism of the metric space R (under the metric induced by the discrete valuation v_P) into the formal power series ring $E[[t]]$ (under the metric induced by the order valuation v_t, giving the lowest-degree term with a nonzero coefficient).
    - For arbitrary rational $a \in K^\times$ with $v_P(a) = d$ we may apply this method to obtain a power series expansion for $t^{-d}a$, and then scale by t^d to obtain a Laurent expansion.
    - Finally, we can define the residue $\mathrm{res}_{P,t}(a)$ of an element $a \in K^\times$ at a prime P with respect to the uniformizer t to be the coefficient a_{-1} in its corresponding Laurent expansion. When P has degree 1, in particular, the residue field R/P is (naturally isomorphic to) K, and we can correspondingly view the residue as an element of K.
- However, giving a reasonable analogue of a differential is more challenging. The most standard approach is to use Weil differentials, but the ideas can also be formulated using the more natural notion of derivations on modules.

- Definition: Let K/F be a function field and M be an F-module (i.e., a vector space over F). A derivation of K/F into M is an F-linear transformation $D : K \to M$ such that $D(ab) = aD(b) + bD(a)$ for all $a, b \in K$.
  - In other words, a derivation is an F-linear function that also obeys the Leibniz formula for the product rule.
  - Example: The usual derivative map $D : F(t) \to F(t)$ with $D(f(t)) = f'(t)$ is a derivation.
  - Exercise: Suppose $D : K \to M$ is a derivation of K/F into M. Show the following:
    1. $D(c) = 0$ for all $c \in F$.
    2. (Quotient Rule) $D(a/b) = [bD(a) - aD(b)]/b^2$ for all $a \in K$, $b \in K^\times$.
    3. (Chain Rule 1) For any $f(x) \in F[x]$ and any $a \in K$, we have $D(f(a)) = f'(a)D(a)$ where f' is the usual formal derivative of f.
  - For any $x \in K \backslash F$, the chain rule shows that the value of a derivation D on $F[x]$ is completely determined by the value $D(x)$, and the quotient rule then extends this observation from $F[x]$ to $F(x)$. If $K = F(x)$ is purely transcendental, then clearly the value of the derivation on K is completely determined by $D(x)$, but otherwise, K is a finite extension of $F(x)$.
  - Unfortunately, if we do not make a good choice for x, then the value $D(x)$ need not determine the value of D on all of K, since for example if $\mathrm{char}(K) = p$ and $x = y^p$ for some $y \in K$ then $D(x) = D(y^p) = 0$, but D need not be identically zero on all of K.
  - Fortunately, this is essentially the only possible problem: one can show in fact that if x is a separating element of K/F (one for which $K/F(x)$ is finite and separable) then the value of $D(x)$ completely determines D on all of K.
  - Exercise: Suppose K/F is a function field with $x \in K \backslash F$ and where $K/F(x)$ is separable. If D_1 and D_2 are derivations from K/F to M and $D_1(x) = D_2(x)$, show that $D_1(a) = D_2(a)$ for all $a \in K$. [Hint: First show D_1 and D_2 agree on $F(x)$. Then for any $y \in K$, apply D_1 and D_2 to its minimal polynomial $m(y) = 0$ over $F(x)$; separability ensures that $m'(y)$ is not the zero polynomial.]
  - It can likewise be shown that if x is a separating element of K/F, then a derivation defined on $F(x)/F$ extends uniquely to a derivation defined on K/F. (We omit the precise details, since they are fairly technical and unenlightening.)
  - Applying this extension result to the natural differentiation map $D : F(x) \to K$ with $D(f(x)) = f'(x)$ yields that for any separating element x of K/F there exists a derivation $D_x : K \to K$ such that $D_x(x) = 1$: this map corresponds to "derivation with respect to x".
  - Exercise: For a function field K/F, let Der_K denote the space of derivations $D : K \to K$.
    1. Show that Der_K is a K-vector space under pointwise addition and scalar multiplication.
    2. Show that for any derivation $D \in \mathrm{Der}_K$ and any $a \in K^\times$ we have $D = D(x) \cdot D_x$. Deduce that Der_K is 1-dimensional (as a K-vector space).

- We can finally define differentials in terms of derivations, as follows:

- Definition: Let K/F be a function field and Z be the set of ordered pairs (a, x) such that $a \in K$ and x is a separating element of K/F. Define an equivalence relation $\sim$ on pairs via $(a, x) \sim (b, y)$ when $b = a \cdot D_y(x)$, and denote the equivalence class of (a, x) as the differential $a\,dx$.
  - Exercise: If x and y are separating elements of a function field K/F, show the chain rule: that $D_x = D_x(y) \cdot D_y$ as functions on K. Deduce that $\sim$ is in fact an equivalence relation.
  - We have a natural K-module structure on the space Diff_K of differentials of K/F, as follows: for any fixed separating element $z \in K$ and any differentials $a\,dx$ and $b\,dy$, we have $a\,dx = [aD_z(x)]\,dz$ and $b\,dy = [bD_z(y)]\,dz$ by the chain rule, so we may define $a\,dx + b\,dy = [aD_z(x) + bD_z(y)]\,dz$.
  - It is easy to see by another application of the chain rule that this addition operation is well defined, as is the natural scalar multiplication $c \cdot (a\,dx) = (ca)\,dx$ for $c \in K$.
  - Exercise: Show that the map $d : K \to \mathrm{Diff}_K$ via $d(a) = 1\,da$ when a is a separating element and $d(a) = 0$ when a is non-separating is a derivation of K/F.

- The space Diff_K is the natural "module of differentials" of K/F, in that any other derivation of K/F into a module M must factor through Diff_K: explicitly, if $D : K \to M$ is a derivation, then there exists a unique F-linear map $\mu : \mathrm{Diff}_K \to M$ with $D = \mu \circ d$.
- The module of differentials Diff_K is a 1-dimensional K-module with basis $\{dz\}$ for any separating $z \in K$, as follows from the observations above: the elements of Diff_K are $a\,dz$ for $a \in K$, and $1\,dz \neq 0$ since $(1, z)$ is not equivalent to $(0, z)$.
- Since this module is 1-dimensional, we may then define the quotient of one differential by another by setting $(a\,dz)/(b\,dz) = a/b \in K$. In particular, if $x \in K$ is separating, then the quotient dy/dx is defined for any $y \in K$, and indeed we have $dy/dx = D_x(y)$ by the chain rule. (This also explains why we call it the chain rule, since in this formulation $D_x(z) = D_x(y) \cdot D_y(z)$ says that $dz/dx = (dy/dx)(dz/dy)$.)
- The differential quotient also shows up when calculating residues with respect to different uniformizers: if s and t are both uniformizers at a prime P of degree 1 (so that the residue of a is an element of K), then $\mathrm{res}_{P,s}(a) = \mathrm{res}_{P,t}(a \cdot ds/dt)$, as one can show by comparing the series expansions using the chain rule.
- We may, at last, use this to define the residue of a differential at a prime of degree 1: for any choice of uniformizer t, we set $\mathrm{res}_P(a\,dt) = \mathrm{res}_{P,t}(a)$; the change-of-uniformizer formula above ensures that the residue does not depend on which uniformizer t we select.
- With some effort, one can then prove that for a nonzero differential $d\omega$, the residue $\mathrm{res}_P(d\omega) = 0$ for all but finitely many primes P (i.e., any nonzero differential has finitely many zeroes and poles), and thus we may attach a divisor to a differential $d\omega$ by setting $\mathrm{div}(d\omega) = \sum_P \mathrm{res}_P(d\omega)P$.
- Exercise: If K/F is a function field and F is algebraically closed, show that all nonzero differentials lie in the same divisor class.
- Finally, we obtain an analogue of the residue theorem we observed earlier for $\mathbb{C}$: for any function field K/F where F is algebraically closed and any differential $d\omega$ of K/F, we have $\sum_P \mathrm{res}_P(d\omega) = 0$.
- As a final remark, we note that we can drop the hypothesis that F is algebraically closed by using the Galois action to work over appropriate subfields of F (the effect being that a degree-d prime of F will correspond to a sum of d degree-1 primes over the algebraic closure of F).

Well, you're at the end of my handout. Hope it was helpful.