18.706 - Noncommutative Algebra

Serina Hu, Vasily Krylov, Roman Bezrukavnikov, Merrick Cai

Spring 2023

Contents

1 February 7 - Introduction, simple and semisimple modules, skew fields

Noncommutative algebra studies algebraic phenomena that arise in a variety of contexts in mathematics and physics, wherever one encounters a multiplication rule where the commutativity law $ab = ba$ fails. An example familiar from linear algebra is multiplication of matrices. Noncommutative groups and Lie algebras also come with such a multiplication; we will also require an addition law compatible with multiplication via the distributive law, groups and Lie algebras can be fit into that framework by passing to the group ring and enveloping algebra respectively.

Some approaches to noncommutative algebras are inspired by known results about commutative ones, where we have familiar concepts of radical, localization etc. We will see noncommutative analogues of these concepts later in the course. Another way to relate the noncommutative and the commutative settings is by deforming a commutative multiplication to obtain a noncommutative one; due to its relation to quantum physics this procedure is sometimes called quantization.

Just as in the case of groups or commutative algebras, much of the work with the abstract composition rule comprising the structure of a group or a ring involves realizing it as composition of actual symmetries of a specific set or abelian group, this leads to the concept of an action of a group on a set, and an action of a ring on a module. A noncommutative ring will be the main protagonist of our story, the plot develops as the protagonist acts (on a module)!

The language of categories and functors is ubiquitous in modern algebra, including study of noncommutative rings and modules over them. Its general concepts and their application to rings and modules will be discussed in the lectures.

Powerful tools for study of rings and modules come from homological algebra, we will introduce its basic concepts in the course.

We will look into core topics in noncommutative ring theory such as polynomial identities and rate of growth of algebras, and also touch upon connections of noncommutative algebra to other areas such as number theory (Brauer groups), Lie theory (Amitsur-Levitzki Theorem, Goldie rank) etc.

The course ends with a brief discussion of noncommutative geometry, an area that grew out of an attempt to connect noncommutative algebra to geometry inspired by the success of algebraic geometry which provides such a connection for commutative algebra.

1.1 Rings, Modules, Ideals

Passing to formal math, the main object of study for us will be associative, possibly noncommutative rings.

Definition 1.1: A **ring** $(R, +)$ is an abelian group with a multiplication that is associative and distributes over addition. Unless stated otherwise, rings will be unital (have a multiplicative identity). Homomorphisms of rings are required to send 1 to 1.

Remark 1.2: Associativity is equivalent to the fact that left multiplication commutes with right multiplication.

Definition 1.3: Let R be a ring. The **opposite ring** R^{op} has the same underlying abelian group as R, but left multiplication by a in R^{op} is defined as right multiplication by a in R, i.e.

$$a \cdot_{\text{op}} b = ba.$$

It is clear that $(R^{\text{op}})^{\text{op}} = R$.

Example 1.4: Fields and skew fields are rings. Recall that a skew field (also known as a division ring) is a ring where every nonzero element is invertible.

Example 1.5: Let R be a ring. Then the set of $n \times n$ matrices with entries in R, with matrix addition and multiplication, is also a ring, denoted $\mathrm{Mat}_n(R)$.

Definition 1.6: A **(left) module** M over a ring R is an abelian group equipped with a ring homomorphism $R \to \mathrm{End}(M)$. Equivalently, we have a bilinear map $R \times M \to M$ satisfying $r_1(r_2(m)) = (r_1 r_2)(m)$.
A **submodule** N of M is a subgroup of M closed under the action of R. Given such $N \subset M$, we can also equip M/N with the structure of an R-module.

Example 1.7: If R is a field, then R-modules are vector spaces.

Definition 1.8: A **bimodule** over R is a module with compatible R-module and R^{op}-module structures, i.e. the actions of R and R^{op} commute.

Example 1.9: R is an R-bimodule; the R-module structure is left multiplication and the R^{op}-module structure is right multiplication, and the associativity of multiplication in R implies that these are compatible.

Definition 1.10: A **left ideal** of R is an R-submodule of R. A **right ideal** of R is an R^{op}-submodule of R (treated as an R^{op}-module). A **two-sided ideal** of R is a subbimodule of R.

Remark 1.11: If I is a left ideal, as described in Definition 1.6, R/I is an R-module, and likewise, if I is a right ideal, R/I is an R^{op}-module. If I is a two-sided ideal, then the multiplication of elements in R/I is well-defined, and R/I is a ring.

Definition 1.12: An R-module M is **free** if it is isomorphic to $\bigoplus_{i \in I} R$, where I is some (possibly infinite) index set. If $M \cong R^n$, we say that M has **rank** n. Note that rank is not well-defined in general!

Example 1.13: Every module over a skew field is free. (See linear algebra.)

Remark 1.14: Remember that in the finite case, direct products and direct sums are the same, but in the infinite case, they are not. In an infinite direct sum, all but finitely many elements must be 0.

1.2 Invariant Basis Number Property

Definition 1.15: A ring R has the **invariant basis number (IBN) property** if free modules of different ranks are not isomorphic. That is, rank is well-defined.

Example 1.16: Linear algebra tells us that modules over a skew field satisfy the IBN.

Lemma 1.17: If $\varphi\colon R \to S$ is a ring homomorphism and S satisfies IBN, then so does R.

Proof. To simplify the discussion, let's focus on finite rank modules. Then $\mathrm{Hom}_R(R^n, R^m) = \mathrm{Mat}_{n,m}(R^{\mathrm{op}})$ ($\mathrm{End}_R(R) = R^{\mathrm{op}}$ because any map $R \to R$ commutes with left multiplication, hence is defined by its value at 1, and this can be extended to R^n). If R doesn't satisfy IBN, there exist non-square matrices $A \in \mathrm{Mat}_{n,m}(R^{\mathrm{op}})$, $B \in \mathrm{Mat}_{m,n}(R^{\mathrm{op}})$ so that $AB = 1_m, BA = 1_n$. But applying φ, we then see that $\varphi(A), \varphi(B)$ give an isomorphism between S^n and S^m, contradiction. ◻

Corollary 1.18: Any ring admitting a homomorphism into a skew field satisfies IBN.

Example 1.19: By Zorn's lemma, every commutative ring R has a maximal ideal $\mathfrak{m}$. Then $R \twoheadrightarrow R/\mathfrak{m}$, which is a field, so R has the IBN.

Example 1.20: We will see later that every left Noetherian ring maps to $\mathrm{Mat}_n(D)$ for some n, D a skew field, so it satisfies IBN.

Example 1.21: Let $V = \mathbb{C}^\infty = \bigoplus_{i=1}^\infty \mathbb{C}$. Then $R := \mathrm{End}(V)$ doesn't satisfy IBN. Choose subspaces V_1, V_2 such that $V = V_1 \oplus V_2$ and $V \cong V_1, V_2$. Then consider the ideals $I_i := \{r \mid r|_{V_i} = 0\}$. $R = I_1 \oplus I_2$, but also $R \cong I_1, I_2$.

Corollary 1.22: $R = \mathrm{End}(\mathbb{C}^\infty)$ does not admit a homomorphism into a skew field.

1.3 Simple modules, Schur Lemma

Theorem 1.23: Suppose that every R-module is free. Then R is a skew field.

To prove this, we will use the Schur Lemma about simple modules.

Definition 1.24: A module M is **simple** or **irreducible** if $M \neq 0$ and it has no nontrivial proper submodules.

Example 1.25: R is simple over itself iff R is a skew field. (If R is simple over itself, then R has no nontrivial ideals, so every nonzero element must be invertible.)

Lemma 1.26 (Schur): If M is simple, then $\mathrm{End}_R(M)$ is a division ring.

Proof. Suppose $\varphi : M \to M$ is nonzero. Then $\ker\varphi \neq M$, but M is simple, so $\ker\varphi = 0$. Hence φ is injective. Likewise, $\mathrm{im}\,\varphi \neq 0$, so $\mathrm{im}\,\varphi = M$ and φ is surjective. Thus φ is invertible. □

Corollary 1.27: Any nonzero map of simple modules is an isomorphism. In particular, if M, N are non-isomorphic simple modules, $\mathrm{Hom}(M, N) = 0$.

Lemma 1.28:

a) Every nonzero ring has a simple module.
b) Every proper left ideal in a nonzero ring is contained in a maximal ideal.
c) A proper submodule N in a module M is maximal iff M/N is simple.

Proof. a) will follow from b) and c) because maximal left ideals of R are maximal R-submodules of R. c) is true because the submodules of M/N are in bijection with the submodules of M containing N.
b) follows from Zorn's Lemma. Its conditions are satisfied because for a nested collection $M_0 \subset M_1 \subset \cdots \subset$ of proper submodules in a finitely generated M, $\bigcup M_i = M$ iff some $M_i = M$. □

Remark 1.29: Part b) is also true for finitely generated modules. If M is not finitely generated, b) may not be true. For example, let $R = \mathbb{Z}, M = \mathbb{Q}$. Then M has no maximal proper submodule because you can find a nested collection of submodules of M whose union is also M.

Corollary 1.30: Every finitely generated module has an irreducible quotient.

Proof (of Theorem 1.23). Let L be a simple R-module (that exists by Lemma 1.28 *a*)). It doesn't contain any submodule isomorphic to R^2 because every nonzero element of L generates L. So if L is free, it must be isomorphic to R. But then $\mathrm{End}_R(L) \cong \mathrm{End}_R(R) = R^{\mathrm{op}}$, and $\mathrm{End}_R(L)$ is a skew field by Lemma 1.26. □

1.4 Semisimple modules

Definition 1.31: A module is **semisimple** if it's isomorphic to a direct sum of simple modules.

Example 1.32: $\mathbb{C}[t]/(t^2)$ is not semisimple as a module over itself. However, we do have an exact sequence of $\mathbb{C}[t]/(t^2)$-modules:

$$0 \to \mathbb{C}[t]/(t) \to \mathbb{C}[t]/(t^2) \to \mathbb{C}[t]/(t) \to 0.$$

Lemma 1.33: Let $M = \bigoplus_{i \in I} M_i$ be a semisimple module, M_i are simple modules. Then any submodule $N \subset M$ has a direct complement of the form $\bigoplus_{i \in J} M_i$ for some $J \subset I$.

Proof. Define $S_J := \bigoplus_{i \in J} M_i$ for $J \subset I$. Consider $J \subset I$ such that $S_J \cap N = 0$; check that the union of a nested collection of these J is a subset J' with $S_{J'} \cap N = 0$. Then there exists a maximal such J. $S_J \cap N = 0$ by construction, and $S_J + N = M$. If not, there exists $M_i \not\subset S_J + N$, and we could then replace J with $J \cup \{i\}$, contradiction. □

Theorem 1.34: Every R-module is semisimple iff $R = \prod_{i=1}^{n} \mathrm{Mat}_{n_i}(D_i)$ where the D_i are skew fields.

2 Semisimple modules, socles, Artinian rings, Wedderburn's Theorem

2.1 More on semisimple modules

Example 2.1: Let D be a skew field. Then D^n is a simple module over $\mathrm{Mat}_n(D)$: given any nonzero vector $v \in D^n$, there's a change of basis matrix M such that $Mv = (1, 0, \ldots, 0)$, and we can then use permutation matrices to get all the other basis vectors. Therefore, $\mathrm{Mat}_n(D)(v) = D^n$.

Corollary 2.2: Subquotients and sums of semisimple modules are semisimple.

Proof. First, we show that submodules of semisimple modules are semisimple. Let $M \cong \bigoplus_{i \in I} L_i$ and $N \subset M$ a submodule. Then by Lemma 1.33, $N \oplus \bigoplus_{i \in J} L_i \cong M$. Therefore, the composition

$$N \hookrightarrow N \oplus \bigoplus_{i \in J} L_i \cong M \twoheadrightarrow \bigoplus_{i \in J \setminus I} L_i$$

is an isomorphism and N is semisimple.
Then quotients of semisimple M are of the form M/N for N a submodule, so by the above $M/N \cong \bigoplus_{i \in J} L_i$ and is semisimple.
Finally, $\sum M_i$ is semisimple because there is a surjection $\bigoplus M_i \twoheadrightarrow \sum M_i$, so $\sum M_i$ is a quotient of the semisimple module $\bigoplus M_i$. □

Example 2.3: $\mathrm{Mat}_n(D)$ is semisimple over itself. It can be decomposed as $\bigoplus_{i=1}^{n} \mathrm{Mat}_n(D)(e_i)$ where e_i are the standard basis vectors: each summand is matrices that have zeroes everywhere except the ith column. Therefore, $\mathrm{Mat}_n(D)(e_i) \cong D^n$; combined with Example 2.1, $\mathrm{Mat}_n(D)$ is then semisimple.

2.2 Socles

Definition 2.4: The **socle** of a module M, denoted $\mathrm{Soc}(M)$, is the sum of all semisimple (or simple) submodules of M. Equivalently, it is the maximal semisimple submodule of M.

Example 2.5: Let $M = \mathbb{C}[t]$ as a $\mathbb{C}[t]$-module. Then $\mathrm{Soc}(M) = 0$. Submodules of M are ideals in $\mathbb{C}[t]$, and an ideal is simple iff it contains no other ideals. But if $I \neq 0$, $tI \subsetneq I$, so (0) is the only simple submodule of M.

Example 2.6: Let $M = \mathbb{C}[t]/t^n$ as a $\mathbb{C}[t]$-module. Then $\mathrm{Soc}(M) = t^{n-1}M$ and is one-dimensional. The submodules of M are all of the form $t^m M$, so they are simple iff $m = n-1$; otherwise, $t(t^m M) \subsetneq t^m M$. Hence the only simple submodule of M is $t^{n-1}M$.

Example 2.7: Let G be a finite p-group and k be a field of characteristic p. Let $M = k[G]$ as a $k[G]$-module. Then $\mathrm{Soc}(M) = k$. To see that, we will show that the only simple G-module is k. We will induct on the order of G. Our base case is $G = \mathbb{Z}/p\mathbb{Z}$. Let V be a simple G-module. Because $(\sigma - 1)^p = 0$ for all $\sigma \in G$, $\ker(\sigma - 1) \neq 0 \Rightarrow \ker(\sigma - 1) = V$. So $\sigma = 1$ and V must be the trivial representation.
Now suppose G is an arbitrary p-group and V an irreducible G-module. Then G has a nontrivial center (can be shown by using the class equation), and the center must contain $\mathbb{Z}/p\mathbb{Z}$. In particular $\mathbb{Z}/p\mathbb{Z}$ is a normal subgroup of G, so $V^{\mathbb{Z}/p\mathbb{Z}}$ is a nonzero $G/(\mathbb{Z}/p\mathbb{Z})$-representation. By induction, it contains a copy of the trivial representation, and so V has a G-invariant vector. So $0 \neq V^G \subset V$ and V must be trivial.

2.3 Isotypic components

For a semisimple module $M \cong \bigoplus_i L_i$, the direct sum decomposition is not canonical; for example, vector spaces have many different bases. But we see that the multiplicity of each L_i is fixed: the number of summands L_i isomorphic to L is $\dim_D(\mathrm{Hom}(L, M)), D = \mathrm{End}(L)^{\mathrm{op}}$. Moreover, the sum of such L_i is well-defined because it is generated by the images of all maps $L \longrightarrow M$ (in fact, all embeddings $L \hookrightarrow M$).

Definition 2.8: Using the above notation, **the L-isotypic component** of M is the sum of the images of all embeddings $L \hookrightarrow M$. Equivalently, if $M \cong \bigoplus L_i$, it is $\bigoplus_{L_i \cong L} L_i$.

Proposition 2.9: M is semisimple iff any short exact sequence $0 \longrightarrow M_1 \longrightarrow M \longrightarrow M_2 \longrightarrow 0$ splits.

Proof. If M is semisimple, Lemma 1.33 and Corollary 2.2 imply that every short exact sequence of the above form splits.
So suppose that every short exact sequence of the above form splits. Consider the short exact sequence $0 \longrightarrow \mathrm{Soc}(M) \longrightarrow M \longrightarrow N \longrightarrow 0$; thus we can write $M = \mathrm{Soc}(M) \oplus N$ and the module N has no simple submodules. Notice that any submodule of $N' \subset N$ is a summand of N: consider the complement of $N' + \mathrm{Soc}(M)$ in M and project down to N. Now take some $a \neq 0, a \in N$ and let $N' := Ra \subset N$. By Corollary 1.30, N' has a simple quotient, say L, and by the same argument L must be a summand of N. But N has no simple submodules, a contradiction. ◻

2.4 Classification of semisimple rings

Theorem 2.10: Every R-module is semisimple iff R is semisimple over itself iff $R = \prod_{i=1}^n \mathrm{Mat}_{n_i}(D_i)$ where the D_i are skew fields. (This is an augmented version of Theorem 1.34.)

This is very often applied to R/J, where J is the Jacobson radical of R; the quotient R/J is semisimple and has the same simple modules as R.

Proof. The first equivalence comes from the fact that every R-module is a quotient of a free module, so if R is semisimple, so is R^I, and so are any quotients of R^I (see Corollary 2.2).
If R is a finite product of matrix rings, Example 2.3 implies that R is semisimple over itself.
To show the last implication, assume R is semisimple over itself and write $R = \bigoplus L_i$. This sum is finite because R is cyclic (it is generated by 1), so if the sum were over an index set $i \in I$, we could write $1 = \sum_{i \in J \subset I} l_i$ where $|J| < \infty$ and $l_i \in L_i$, so $R = \bigoplus_{i \in J} L_i$. (The same argument would work for any finitely generated module.) Anyway, write R as the sum of its isotypic components, say

$$\bigoplus_{j \in J} L_j^{d_j}, L_j \neq L_{j'} \Leftrightarrow j \neq j'.$$

We know that

$$R^{\mathrm{op}} = \mathrm{End}_R(R) = \mathrm{End}_R\left(\bigoplus_{j\in J} L_j^{d_j}\right) = \prod_{j\in J} \mathrm{Mat}_{d_j}(\mathrm{End}_R(L_j))$$

and if we let $D_j = (\mathrm{End}_R(L_j))^{\mathrm{op}}$, we get an isomorphism

$$R \cong \prod_{j\in J} \mathrm{Mat}_{d_j}(D_j).$$

□

Remark 2.11: It would seem natural to call rings R semisimple over themselves semisimple. However, there is a separate notion of a simple ring, and not all simple rings are semisimple over themselves (see Example 2.13 below).

2.5 Simple rings and Wedderburn's Theorem

Definition 2.12: A ring R is **simple** if R has no 2-sided ideals except for 0 and R.

Example 2.13: $R = \mathbb{C}\langle x, \partial_x\rangle$ is simple but not semisimple. To see that R is not semisimple, consider $R/R(x\partial_x)$. This module has a surjection to $R/R(\partial_x)$ that does not split (exercise).

Definition 2.14: A ring R is **left (resp. right) Noetherian** if every ascending chain of left (resp. right) ideals of R stabilizes (called the ascending chain condition). Equivalently, every left (resp. right) ideal is finitely generated.

Definition 2.15: A ring R is **left (resp. right) Artinian** if every descending chain of left (resp. right) ideals of R stabilizes (the descending chain condition).

Warning 2.16: Being left Artinian/Noetherian is not equivalent to being right Artinian/Noetherian!

Theorem 2.17 (Wedderburn): Let R be a ring. TFAE:

a) R is simple and (either left or right) Artinian,

b) every R-module is semisimple and R has a unique simple module up to isomorphism,

c) $R \cong \mathrm{Mat}_n(D)$ where D is a skew field.

Proof. The equivalence of *b*) and *c*) follows from Theorem 2.10: if R is a finite product of matrix rings over skew fields, check that $\mathrm{Mat}_n(D)$ is simple over itself, and so R has a unique simple module iff the product only contains one matrix ring. This also shows that *c*) implies *a*).

So suppose that R is left Artinian and simple. Then R has a minimal left ideal (because any descending chain of left ideals will stabilize), call it L. Notice that $LR = \sum_{x\in R} Lx$ is a nonzero two-sided ideal, hence all of R, and $R = LR$. So R as a left R-module is a quotient of $\bigoplus_{x\in R} L$, and R is semisimple over itself. Thus *a*) implies *b*) by use Theorem 2.10. □

3.1 $k[G]$-modules

Example 3.1: Let G be a finite group and k a field of characteristic not dividing $|G|$ (for simplicity, let's say $\operatorname{char} k = 0$, but the result holds in general). Then all $k[G]$-modules are semisimple.
We will show that every short exact sequence $0 \to N \to M \to L \to 0$ splits. WLOG, we can assume that L is finite-dimensional. Tensoring with L^* and using the fact that $\operatorname{Hom}_G(V, W) = (V^* \otimes W)^G$ when V is finite-dimensional, it suffices to show that for $M \twoheadrightarrow L$, the restriction to $M^G \to L^G$ is also onto. But this is true because given $v \in L^G$, choose any preimage of v in M, say $\tilde{v}$, and consider $\frac{1}{|G|} \sum g(\tilde{v})$, which lies in M^G and maps to v.

Corollary 3.2: Suppose that k is algebraically closed and $\operatorname{char} k$ does not divide $|G|$. Then $|G| = \sum(\dim \rho_i)^2$ where the ρ_i are the isomorphism classes of simple $k[G]$-modules.

Proof. The only finite skew field extensions of k are trivial if k is algebraically closed. Hence, by Theorem 2.10 $k[G]$ semisimple means it can be written as $\prod_{i=1}^n \operatorname{Mat}_{d_i}(k)$, and the simple $k[G]$-modules are exactly k^{d_i}, while the dimension of $k[G]$ over k is $\sum d_i^2$. □

3.2 Density Theorem

Theorem 3.3 (Density Theorem): Let L be a simple R-module and $D = \operatorname{End}_R(L)$. Then given any finite set $x_1, \dots, x_n, y_1, \dots, y_n \in L$ with the x_i linearly independent over D, there exists $r \in R$ such that $r(x_i) = y_i$.

Proof. Let $\boldsymbol{x} = (x_1, \dots, x_n)$. We want to show that the map $R \to L^n$ taking $r \mapsto r\boldsymbol{x}$ is onto. Suppose that $R\boldsymbol{x} \subset L^n$ is a proper submodule, say N. Since L^n is semisimple, we can then decompose $L^n = N \oplus S$, $S \neq 0$. Therefore, $D^n = \operatorname{Hom}_R(L, L^n) = \operatorname{Hom}_R(L, N) \oplus \operatorname{Hom}_R(L, S)$. Therefore, there exists some $(d_1, \dots, d_n) \in D^n$ annihilating the proper subspace $\operatorname{Hom}(L, N)$ (acting via the dot product), so the x_i are linearly dependent, a contradiction. □

Remark 3.4: The submodules in an isotypic component $L^n \subset M$ are in bijection with vector subspaces in $D^n, D = \operatorname{End}(L)$. The correspondence sends $N \subset L^n$ to $\operatorname{Hom}(L, N) \subset \operatorname{Hom}(L, L^n) = D^n$ (exercise).

Corollary 3.5: If L is finite-dimensional simple over $D := \operatorname{End}_R(L)$, then there is a surjection $R \twoheadrightarrow \operatorname{End}_D(L) \cong \operatorname{Mat}_n(D), n = \dim_D(L)$.

Example 3.6: This is not true if M is infinite-dimensional over D. For example, let $R = \operatorname{End}(\mathbb{C}^\infty)$ and $M = \mathbb{C}^\infty$. Then $D = \operatorname{End}_R(M) = \mathbb{C}$ but there is no surjection $R \to \operatorname{End}_D(M)$ (see Corollary 1.22).

3.3 Noetherian and Artinian modules

Definition 3.7: A module is **Noetherian (resp. Artinian)** if every ascending (resp. descending) chain of submodules stabilizes.

Remark 3.8: We'll see that every Artinian ring is also Noetherian, but this is not true for modules.

Example 3.9: Let $R = \mathbb{Z}$. Then $M = \mathbb{Z}$ is a Noetherian module, but it is not Artinian because $(p) \supset (p^2) \supset (p^3) \supset \cdots$ is an infinite descending chain of submodules. Meanwhile, $N = \mathbb{Q}/\mathbb{Z}$ is Artinian, but it is not Noetherian, because $\frac{1}{p}N \subset \frac{1}{p^2}N \subset \cdots$.

Proposition 3.10: A module is Noetherian iff every submodule is finitely generated.

Proof. Let M be an R-module. If every $N \subset M$ is finitely generated, suppose we had an ascending chain of submodules $M_1 \subset M_2 \subset \cdots \subset$ and consider $N = \bigcup M_i$. Because N is finitely generated, say with generators $x_1, \ldots, x_d$, there exists some i with $M_i \supset \{x_1, \ldots, x_d\}$, and the ascending chain stabilizes at M_i.
Now suppose that M is Noetherian and $N \subset M$ is a submodule. Obtain a list of generators $x_i \in N$ by taking $x_1 \neq 0$ and x_i any element not in $N_{i-1} := \langle x_1, \ldots, x_{i-1} \rangle$. The ascending chain $N_1 \subset N_2 \subset \cdots$ must stabilize eventually, say at N_d, and $x_1, \ldots, x_d$ then generate N. □

Proposition 3.11: If $0 \to M_1 \to M \to M_2 \to 0$ is a short exact sequence and M_1, M_2 are Noetherian (resp. Artinian), then M is also Noetherian (resp. Artinian).

Proof. Clear. □

3.4 Composition Series

Definition 3.12: A **composition series** of a module M is a filtration $M_0 = 0 \subsetneq M_1 \subsetneq M_2 \subsetneq \cdots \subsetneq M_n = M$ where M_i/M_{i-1} is simple for all i. That is, the filtration has simple associated graded subquotients. If M has a composition series, we say that it is of **finite length** and say that M has **length** n.

Lemma 3.13: A module M has finite length iff M is both Noetherian and Artinian.

Proof. First, suppose M has a composition series. Then induct on the length of M. If M has length 1, it's simple, and therefore both Noetherian and Artinian. If M has length n, then $0 \to M_{n-1} \to M \to L \to 0$ and M_{n-1}, L are Noetherian and Artinian by induction, so M also is.
Now suppose M is both Noetherian and Artinian. Because M is Artinian, by Zorn's Lemma any nonempty collection of submodules has a minimal element. So let $M_1 \subset M$ be a minimal nonzero submodule; it must be a simple submodule. Now inductively define M_{i+1} to be the minimal submodule properly containing M_i; this will exist unless $M_i = M$, and M_{i+1}/M_i will be simple. This chain of submodules will terminate because M is Noetherian, so $M_n = M$ for some n and we have constructed a composition series for M. □

Definition 3.14: Let $M_1 \subset \cdots \subset M_n = M$ be a composition series for M. The **associated graded** of the composition series is

$$\operatorname{gr}(M) := \bigoplus_{i=1}^{n} M_i/M_{i-1}.$$

Theorem 3.15 (Jordan-Hölder): Given two composition series M_i, M_i' of M, $\operatorname{gr}(M) = \operatorname{gr}'(M)$. Equivalently, the number of irreducible subquotients isomorphic to a given simple module L is independent of the choice of filtration.

Proof. Induct on the length of M_i. If M_i has length 1, M is simple and both filtrations contain only M with multiplicity 1. If not, consider the smallest j such that $L = M_1 \subset M_j'$. Since $L \not\subset M_{j-1}'$, there is a nonzero map $L \to M_j'/M_{j-1}' = \operatorname{gr}_j'(M)$, and a nonzero map between simples is an isomorphism. Hence $\operatorname{gr}_j'(M) \cong L$.
Therefore, M/M_1 has two filtrations: one given by $\bar{M}_i = M_{i+1}/M_1$ and one defined by $\bar{M}_i'$ is the image of M_i' when $i < j$ and M_{i+1}'/M_1 when $i \geqslant j$. We know that we get $\overline{\operatorname{gr}}(M) = \overline{\operatorname{gr}}'(M)$ from removing one copy of L from $\operatorname{gr}(M)$ and $\operatorname{gr}'(M)$, so by induction, $\operatorname{gr}(M) = \operatorname{gr}'(M)$. □

Remark 3.16: Inspecting the proof of Theorem 3.15, we see that a stronger version of it holds. This stronger version claims that for two composition series $0 \subset M_1 \subset \ldots \subset M_a = M$, $0 \subset M_1' \subset \ldots \subset M_b' = M$ of M there exists a *canonical* bijection $\sigma \colon \{1, \ldots, a\} \xrightarrow{\sim} \{1, \ldots, b\}$ and a *canonical* isomorphism $M_i/M_{i-1} \xrightarrow{\sim} M_{\sigma(i)}'/M_{\sigma(i)-1}'$. This version of the theorem is interesting already for $R = k$ (so M is a finite-dimensional vector space): in this case, composition series of M are flags of subspaces in M, and σ describes a "relative position" of these two flags with respect to each other.

Definition 3.17: Let $\mathcal{M}$ be a collection of R-modules closed under subquotients. The **Grothendieck group** $K(\mathcal{M})$ is the free abelian group generated by $[M]$, $M \in \mathcal{M}$, subject to the relations $[M] = [M_1] + [M_2]$ when there is a SES $0 \to M_1 \to M \to M_2 \to 0$.

Remark 3.18: For A an abelian group, any function $\mathcal{M} \to A$ additive on subquotients then induces a map $K(\mathcal{M}) \to A$. For example, if $R = D$ and $\mathcal{M}$ consists of the finite-dimensional vector spaces, dimension is such a function.

Corollary 3.19: Let $\mathcal{M}$ be the modules of finite length over R. Then $K(\mathcal{M})$ is freely generated by $[L]$ for (isomorphism classes of) irreducible modules L.

Proof. The existence of a composition series for each $M \in \mathcal{M}$ means that the $[L]$ generate $K(\mathcal{M})$. To see that the $[L]$ have no relations, notice that Jordan-Hölder implies that there's a well-defined homomorphism $K(\mathcal{M}) \to \mathbb{Z}$ sending $[M]$ to the multiplicity of L in the Jordan-Hölder series of M. Thus every $[M]$ has a unique decomposition into the $[L]$. □

3.5 Jacobson Radical

Definition 3.20: The **Jacobson radical** $J = J(R)$ of a ring R is the intersection of the annihilators of all simple R-modules. In particular, $J(R)$ kills all simple, hence semisimple, modules!

The Jacobson radical has many characterizations.

Lemma 3.21: For $a \in R$ TFAE:

- *a)* $a \in \operatorname{Ann}(L)$ for all simple R-modules L (i.e., $a \in J(R)$),
- *b)* $a \in I$ for all maximal left ideals I,
- *c)* $1 - xa$ has a left inverse for all x,
- *d)* $1 - xay$ has an inverse for all x, y,
- *e)* $1 - ax$ has a right inverse for all x,
- *f)* $a \in I$ for all maximal right ideals I,
- *g)* $a \in \operatorname{Ann}(L)$ for all simple R^{op}-modules L (i.e., $a \in J(R^{\mathrm{op}})$).

4 February 16 - Socle and cosocle filtrations, Jacobson radical, Krull-Schmidt

4.1 Socle and cosocle filtrations

Definition 4.1: The **socle filtration** $M_1 \subset M_2 \subset M_3 \subset \cdots \subset M$ of a module M is defined inductively as follows: M_1 is the socle of M (see Definition 2.4) and M_i is the preimage of the socle of M/M_{i-1} in M.

Remark 4.2: The socle filtration can be generalized to transfinite numbers (e.g. ordinals), in which case it is called the **Loewy filtration**, but we won't talk about it.

Definition 4.3: The **cosocle filtration** $M \supset M^1 \supset M^2 \supset \cdots$ of an Artinian module M is also defined inductively: M^1 is the kernel of the map from M to its maximal semisimple quotient (called the **cosocle**), M^2 is the kernel of the map from M^1 to its cosocle, and so on.

Remark 4.4: If M is Artinian, then the cosocle filtration always exists, but this is not true in general because M may not necessarily have a maximal semisimple quotient. One could consider all possible simple quotients $M \twoheadrightarrow L_i$ and get a map $M \longrightarrow \prod L_i$, but this infinite product need not be semisimple. For example, this occurs when $R = \mathbb{Z}$; then $\prod_p \mathbb{Z}/p\mathbb{Z}$ is not semisimple.
But if M is Artinian, we know the intersection of the kernels of all maps $M \twoheadrightarrow L_i$ is equal to the intersection of the kernels of finitely many such maps: we can order the kernels of all maps $M \longrightarrow \prod_{i=1}^n L_i$ to create a decreasing sequence of submodules, which must stabilize. Hence, there exists a maximal quotient corresponding to the stabilized kernel, $M \longrightarrow \prod_{i=1}^n L_i$. By definition, any map $M \longrightarrow N$ where N is semisimple factors through this image, so $\prod_{i=1}^n L_i$ is the maximal semisimple quotient.

Example 4.5: Let $R = \mathbb{C}[t]$ and suppose that M is a finite-dimensional R-module where t acts nilpotently. Then $M_i = \ker(t^i)$ and $M^i = \operatorname{im}(t^i)$.

Example 4.6: Let $R = \mathbb{Z}/72\mathbb{Z}$ and suppose that M is a free rank 1 R-module, i.e. $M = \mathbb{Z}/72\mathbb{Z}$. Let us compute the socle and cosocle filtrations.
The socle is given by the sum of the simple submodules, which are $36M$ and $24M$, hence $M_1 = 12M$. Then $N := M/M_1 \cong \mathbb{Z}/12\mathbb{Z}$ with the standard action of R. The simple modules N are $4N$ and $6N$, which sum to $2N$, hence M_2 is the preimage in M of $2N$ under the obvious quotient map $M \twoheadrightarrow N$, which is just $2M$. Now $M/2M \cong \mathbb{Z}/2\mathbb{Z}$ which is simple, hence the socle filtration is given by $0 \subset (M_1 = 12M) \subset (M_2 = 2M) \subset M$.
Let us now compute the cosocle filtration. The irreducible quotients are $\mathbb{Z}/2\mathbb{Z}$ and $\mathbb{Z}/3\mathbb{Z}$, hence the cosocle is given by $M^1 = \ker(\mathbb{Z}/72\mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/3\mathbb{Z}) = 6M \cong \mathbb{Z}/12\mathbb{Z}$. Once again, the irreducible quotients of M^1 are $\mathbb{Z}/2\mathbb{Z}$ and $\mathbb{Z}/3\mathbb{Z}$, so $M^2 = 6M^1 = 36M \cong \mathbb{Z}/2\mathbb{Z}$. Finally, M^2 is irreducible so $M^3 = 0$, hence the cosocle filtration is given by $M \supset (M^1 = 6M) \supset (M^2 = 36M) \supset (M^3 = 0)$.

4.2 Jacobson radical cont.

Proof (of Lemma 3.21). a) implies b): for any $\mathfrak{m} \subset R$, $R/\mathfrak{m}$ is simple, so a annihilates $R/\mathfrak{m} \Rightarrow a \in \mathfrak{m}$.
b) implies a): the annihilator of every simple module is a proper ideal in R, thus contained in some maximal left ideal.
c) implies b): if there exists a maximal ideal $\mathfrak{m}$ with $a \notin \mathfrak{m}$, then there exists x such that $xa \equiv 1 (\operatorname{mod} \mathfrak{m})$. Hence $1 - xa$ is not invertible.
b) implies c): first note that $t \in R$ is left invertible iff $Rt = R$ iff t does not belong to a proper left ideal. By Zorn's Lemma, this is equivalent to $t \notin \mathfrak{m}$ for some maximal left ideal. So if $a \in \mathfrak{m}$ for all maximal $\mathfrak{m}$, $1 - xa \notin \mathfrak{m}$ and $1 - xa$ is left invertible.
d) implies c) follows from setting $y = 1$.
c) implies d): the set of all a satisfying a), b), c) forms a 2-sided ideal by a). So xay also lies in this ideal and $1 - xay$ has a left inverse by c); say it is $1 - b$. Then $(1 - b)(1 - xay) = 1$, and so b also lies in the two-sided ideal. By c), $1 - b$ then has a left inverse, which implies that $1 - xay$ is invertible.
Since d) is left-right symmetric, e), f), and g) follow. □

Remark 4.7: If a is nilpotent with $a^n = 0$, then $1 - a$ is invertible with inverse $1 + a + \cdots + a^{n-1}$. Hence, if xay is nilpotent for all x, y, then $a \in J$.

Example 4.8: Let $R = \mathbb{C}[x_1, \ldots, x_n]/I$. The Jacobson radical of R is $\sqrt{I}/I$, which follows from Hilbert's Nullstellensatz.

Example 4.9: If R is a commutative local ring, then $J(R) = \mathfrak{m}$, the unique maximal ideal.

Example 4.10: If $R \subset \operatorname{Mat}_n(k)$ is the subalgebra of upper triangular matrices, then $J(R)$ is the *strictly* upper triangular matrices (zeroes on the diagonal).

4.3 Local rings and indecomposable modules

Definition 4.11: A ring R is **local** if all non-invertible elements form an ideal, in which case said ideal is $J(R)$. If R is local, $R/J(R)$ is a skew field.

Definition 4.12: A module M is **indecomposable** if it cannot be decomposed as a direct sum of nonzero submodules $M_1 \oplus M_2$.

Example 4.13: Let $R = \mathbb{C}[t]$, $M = \mathbb{C}^n$, and t acts by some matrix A. Then M is indecomposable iff A has only one Jordan block.

Remark 4.14: M is indecomposable iff $\mathrm{End}_R(M)$ has no nontrivial idempotents, i.e. elements e such that $e^2 = e$. If $e \in \mathrm{End}_R(M)$, then we could write $M = Me \oplus M(1-e)$: $\ker e = \mathrm{im}(1-e)$ because $(1-e)^2 = 1-e$, so $em = 0 \Leftrightarrow (1-e)m = m \Leftrightarrow (1-e)n = m$ for some n.
Conversely, given a decomposition $M = M_1 \oplus M_2$, we could set $e = \pi_{M_1} : M \twoheadrightarrow M_1$.

Remark 4.15: If we took an idempotent of R instead of $\mathrm{End}_R(M)$, we would still get a splitting $M = eM \oplus (1-e)M$, but this would only be a direct sum of abelian groups, not of R-modules.

Proposition 4.16: If M is indecomposable of finite length, then $\mathrm{End}_R(M)$ is local.

Lemma 4.17: If M is an indecomposable finite length module (therefore Noetherian and Artinian), every $a \in \mathrm{End}_R(M)$ is either nilpotent or invertible.

Proof. For every $a \in \mathrm{End}(M)$, consider the chains $\ker(a) \subset \ker(a^2) \subset \cdots \subset$ and $\mathrm{im}(a) \supset \mathrm{im}(a^2) \supset \cdots \supset$. Because M is finite length, it is both Artinian and Noetherian (see Lemma 3.13), these both stabilize. Let $b = a^n$ where n is such that $\ker(a^{n+1}) = \ker(a^n)$ and $\mathrm{im}(a^{n+1}) = \mathrm{im}(a^n)$. Thus $\ker(b^2) = \ker(b)$, $\mathrm{im}(b^2) = \mathrm{im}(b)$. We claim that then $M = \ker(b) \oplus \mathrm{im}(b)$.
For $x \in \mathrm{End}(M)$, since $\mathrm{im}(b^2) = \mathrm{im}(b)$, there exists y such that $b^2y = bx$. So $x - by \in \ker(b)$ and $x = (x-by)+by \Rightarrow M = \ker(b) + \mathrm{im}(b)$. To see that it's the direct sum, note that $x \in \ker(b) \cap \mathrm{im}(b)$ implies $x = by$ and $bx = b^2y = 0$, but $\ker(b^2) = \ker(b)$, so $by = 0 \Rightarrow x = 0$. Hence $\ker(b) \cap \mathrm{im}(b) = \{0\}$.
Since M is indecomposable, either $\ker(b) = 0$ and $\mathrm{im}(b) = M$ or $\mathrm{im}(b) = 0$ and $\ker(b) = M$. If $\ker(b) = 0$ and $\mathrm{im}(b) = M$, then $\ker(a) = 0$ and $\mathrm{im}(a) = M$ also and so a is invertible. If $\mathrm{im}(b) = 0$, then $b = 0$, so a is nilpotent. □

Proof (of Proposition 4.16). If $a \in \mathrm{End}_R(M)$ is not invertible, it's nilpotent. Hence $\ker(a) \neq 0$. So xa is also not invertible, hence nilpotent. By the same argument, xay is also not invertible, hence nilpotent. By Remark 4.7, $1 - xay$ is invertible for all x, y and thus $a \in J(R)$. So all of the non-invertible elements form the ideal $J(R)$, hence R is local. □

4.4 Krull-Schmidt

Theorem 4.18 (Krull-Schmidt):

a) Every finite length module can be decomposed as a direct sum of indecomposable modules.
b) For any two such decompositions, the multisets of isomorphism classes of the indecomposable summands coincide.

Example 4.19: Let $R = \mathbb{C}[t]$. Then a finite length module is a finite-dimensional vector space and t acts by a matrix. Indecomposable modules correspond to matrices with a single Jordan block, so in this case, Krull-Schmidt is equivalent to saying every matrix has a (essentially unique) Jordan normal form.

Proof (of Theorem 4.18). The proof that such a decomposition exists only requires our module to be either Noetherian or Artinian, but not both. Suppose that M cannot be written as a direct sum of indecomposables. So M is

not indecomposable, which means it has a decomposition $M = M_1 \oplus M_2$ but one of M_1, M_2 is not a direct sum of indecomposables, WLOG M_1. Then we can split M_1, and inductively continue the process indefinitely. This gives us both an infinite descending chain of submodules (the submodules we split at every step) and an infinite ascending chain of submodules (the complement of those submodules), one of which stabilizes, a contradiction.
However, uniqueness requires M to be of finite length.
Let P, Q be any two R-modules. Let $S = \mathrm{End}_R(P)^{\mathrm{op}}$. Then $\mathrm{Hom}_R(P,Q)$ is a left S-module and $\mathrm{Hom}_R(Q,P)$ is a right S-module. Even better, we have a pairing

$$\begin{array}{rcl} \mathrm{Hom}_R(P,Q) \times \mathrm{Hom}_R(Q,P) & \to & \mathrm{Hom}_R(P,P) = S \\ (f,g) & \mapsto & g \circ f. \end{array}$$

If P and Q are indecomposable, then S is local with maximal ideal $\mathfrak{m}_S = J(S)$. Then we claim that the image of this pairing lands in $\mathfrak{m}_S$ iff $P \not\cong Q$. Suppose that there exists f, g with $g \circ f$ invertible; then $Q \cong P \oplus \ker(g)$. This contradicts the indecomposability of Q unless $P \cong Q$.
Now consider $\overline{\mathrm{Hom}_R(P,Q)} := \mathrm{Hom}_R(P,Q)/\mathfrak{m}_S \mathrm{Hom}_R(P,Q)$ and likewise define $\overline{\mathrm{Hom}_R(Q,P)}$. Both of these are modules over the skew field $D_S := S/\mathfrak{m}_S$, i.e. vector spaces, so we get a D-bilinear pairing

$$\overline{\mathrm{Hom}_R(P,Q)} \times \overline{\mathrm{Hom}_R(Q,P)} \to D$$

and this pairing is nonzero iff $P \cong Q$.
Moreover, if Q is not indecomposable, but instead a direct sum $Q_1 \oplus Q_2$, then

$$\overline{\mathrm{Hom}_R(P,Q)} = \overline{\mathrm{Hom}_R(P,Q_1)} \oplus \overline{\mathrm{Hom}_R(P,Q_2)}$$

and likewise for $\overline{\mathrm{Hom}_R(Q,P)}$, and these direct sum decompositions are compatible with the pairing.
Therefore, if $M = \bigoplus_{i=1}^n Q_i$ for Q_i indecomposable, we can likewise decompose $\overline{\mathrm{Hom}_R(P,M)}$ and $\overline{\mathrm{Hom}_R(M,P)}$ and deduce that the number of Q_i isomorphic to a given P is the rank of the pairing $\overline{\mathrm{Hom}_R(P,M)} \times \overline{\mathrm{Hom}_R(Q,M)} \to D$. This is independent of the decomposition, so the multiplicities of the isomorphism classes of the indecomposables are unique. □

5.1 Interlude on quiver representations

While the indecomposables of $\mathbb{C}[t]$ have a nice classification via Jordan normal form, this is generally a wild problem.

For example, one way we can generalize this is by asking how to parametrize finite sets of subspaces of a vector space V. For example, how can we parametrize triples V_1, V_2, V where $V_1, V_2 \subset V$? Say two triples V_1, V_2, V and V_1', V_2', V' are equivalent when there is an isomorphism $V \cong V'$ that sends $V_i \mapsto V_i'$. This is not so bad – these triples are determined up to equivalence by the integers $\dim V_i$, $\dim V$, and $\dim(V_1 \cap V_2)$.

Another nice example is considering invariants for $V_1, V_2, V_3, V_4 \subset V$ when $V = \mathbb{R}^2$. If we require all 4 to be 1-dimensional subspaces of $\mathbb{R}^2$, we want to parametrize quadruples of lines on the plane. In general position, no two coincide, and an isomorphism will take one configuration to the other when they have the same cross ratio. So in this case, our invariant is a general element of $\mathbb{R}$, not a bunch of integers.

More generally, you could ask how to describe any number of subspaces in a vector space. We can rephrase this question in the language of quiver representations. Recall that a quiver is an oriented graph, and a representation of a quiver is just an assignment of a vector space to each vertex and a map between the corresponding vector spaces for each edge. For example, a representation of the below quiver is 4 vector spaces, one for each vertex, and maps between them.

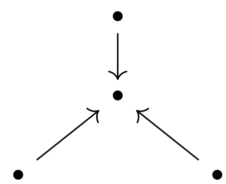

That is, a representation looks like

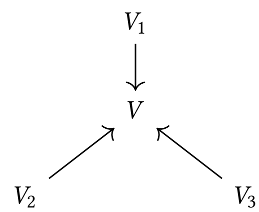

and we can define isomorphisms and direct sums of representations, hence speak about indecomposable and simple representations of this quiver.

Representations of a quiver Q are equivalent to modules of its path algebra $A(Q)$. So Krull-Schmidt tells us that the decomposition of a finite-dimensional representation into indecomposables has unique multiplicities. Fact: the above quiver has 12 indecomposables, so there are 12 invariants necessary to describe a representation of this quiver (one for each indecomposable multiplicity). The dimension of V_1, V_2, V_3 is at most 1, while the dimension of V is at most 2, in each indecomposable. In three of these, the maps aren't injective. So quadruples $V_1, V_2, V_3 \subset V$ are parametrized by 9 invariants, and in fact, we can express these explicitly as intersections.

5.2 Primitive and semi-primitive rings

Definition 5.1: We say a ring R is **semi-primitive** if $J(R) = 0$. Equivalently, $R \hookrightarrow \mathrm{End}(M)$ for some semisimple R-module M. Since $J(R) = J(R^{\mathrm{op}})$, we could also say that $R \hookrightarrow \mathrm{End}(M)$ for a semisimple R^{op}-module M.

Definition 5.2: We say a ring R is **(left, right) primitive** if R has a faithful simple (left, right) R-module, that is, $R \hookrightarrow \mathrm{End}(M)$. There exist rings that are left but not right primitive.

So primitive rings correspond to having a faithful simple module, while semi-primitive rings correspond to having a faithful semisimple module. What distinguishes them is the scenario where every simple module has nontrivial annihilator, so the ring is not primitive, but there exist some simple modules $L_1, \dots, L_n$ whose annihilators intersect to 0, hence the semisimple module $L_1 \oplus L_2 \oplus \cdots \oplus L_n$ has trivial annihilator. In this situation, the ring is semi-primitive but not primitive.

Example 5.3: Simple rings are both left and right primitive: every simple module of a simple ring R has to be faithful, because $\ker(R \to \mathrm{End}(L))$ is a 2-sided ideal in R, hence 0.

Example 5.4: Primitive rings need not be simple. For example, $R = \mathrm{End}(\mathbb{C}^\infty)$ is primitive but not simple. It is primitive because $\mathbb{C}^\infty$ is a simple R-module, but R is not simple because operators of finite rank in R form a two-sided ideal.

Example 5.5: Here's a more "real-life" example. consider $R = U(\mathfrak{sl}_2)/(C)$, where C is the Casimir $ef + fe + \frac{h^2}{2}$ (a central element). We claim this is primitive but not simple. First, R can be identified with the ring S of global differential operators on $\mathbb{P}^1$. Verifying this is a hard exercise; here is an outline:
On each copy of $\mathbb{C}$, the differential operators are generated by $x, \frac{\partial}{\partial x}$. To move between copies, note that $\frac{\partial}{\partial x} = -x^{-2}\frac{\partial}{\partial x^{-1}}$. You can show that the global vector fields on $\mathbb{P}^1$ are generated by $\frac{\partial}{\partial x}, 2x\frac{\partial}{\partial x}, x^2\frac{\partial}{\partial x}$, and the Lie algebra they generate (via taking the commutator of vector fields) has the same relations as $\mathfrak{sl}_2$. This gives us a map $U(\mathfrak{sl}_2) \to S$, and you can check that it kills C, so we get a map $R \to S$, and you then show this map is an isomorphism.
(The geometric explanation for this: SL_2 acts on $\mathbb{P}^1$, so we get a map from $\mathfrak{sl}_2$ to the Lie algebra of vector fields on $\mathbb{P}^1$. This has a far-reaching generalization describing differential operators on a flag variety as an appropriate quotient of the universal enveloping algebra modulo an ideal generated by central elements.)
Anyway, to construct a faithful simple R-module, note that the differential operators on $\mathbb{A}^1$ act on $\mathbb{C}[x]$, and this induces an action of differential operators on $\mathbb{P}^1$ on $\mathbb{C}[x]/\mathbb{C}$. Exercise: verify this is in fact a faithful simple R-module.

So now we have several examples where primitive rings need not be simple. However, if we add the condition that our ring must be Artinian (e.g. a finite-dimensional algebra over a field), then every primitive ring is in fact simple.

Proposition 5.6: A (left or right) Artinian semi-primitive ring has the form $\prod_{i=1}^n \mathrm{Mat}_{n_i}(D_i)$, so (left or right) Artinian primitive rings are of the form $\mathrm{Mat}_n(D)$, hence simple.

Proof. Suppose that R is Artinian and semi-primitive, i.e. $J(R) = 0$. We can also write $J(R) = \bigcap I_\alpha$ where the intersection is over all maximal left ideals I_α. Because R is Artinian, there exists a finite subset of the I_α such that $\bigcap_{i=1}^n I_i = 0$ (consider the infinite descending chain of ideals $I_1, I_1 \cap I_2, \ldots$. This must stabilize, but also $\bigcap I_\alpha = 0$, so it must stabilize at 0).
Therefore, we have an injection $R \hookrightarrow \bigoplus_{i=1}^n R/I_i$. Because each I_i is maximal, R/I_i is simple, so R is a submodule of a semisimple module, hence is semisimple itself. Then by Theorem 2.10, R is a finite product of matrix algebras. Then the second part follows since the simple representations of R will be $\mathrm{Mat}_{n_i}(D_i)$, and these are not faithful unless $R = \mathrm{Mat}_n(D)$. □

Corollary 5.7: Suppose R is Artinian. Then M is semisimple iff $J(R)M = 0$. The socle filtration on M has $M_i = \ker J(R)^i$ and the cosocle filtration is $J(R)^i M$.

Proof. In one direction, if M is semisimple, then by definition $J(R)$ annihilates all simple, hence all semisimple, modules. In the other, suppose that $J(R)$ acts trivially. Then M is a quotient over $R/J(R)$. $J(R/J(R)) = 0$, so this quotient is semi-primitive. It is also Artinian, so M is a module over $\prod_{i=1}^n \mathrm{Mat}_{n_i}(D_i)$. This ring is semisimple, so M is semisimple.
The second statement follows from the first; for example, the socle is the maximal semisimple submodule of M, which must then be the kernel of $J(R)$, and so on. □

Corollary 5.8 (A Version of Nakayama): Suppose M is a finitely generated R-module such that $J(R)M = M$. Then $M = 0$.

Proof. If M is nonzero, we know that M has a simple quotient by Corollary 1.30, call it L, and $J(R)L = 0$. Then $J(R)M \neq M$. □

6 February 28 - Artinian rings are Noetherian, projective covers

6.1 The Akizuki-Hopkins-Levitzki Theorem (Artinian rings are Noetherian)

Lemma 6.1: If R is Artinian, then $J = J(R)$ is a nilpotent ideal, i.e. there exists some $n > 0$ such that $J^n = 0$.

Proof. Saying that $J^n = 0$ is equivalent to saying that $x_1 x_2 \cdots x_n = 0$ for all $x_i \in J$. Consider the decreasing chain $J \supset J^2 \supset \cdots \supset$, which stabilizes because R is Artinian. So let $I = J^n = J^{n+1}$; then $I = I^2$ also. If $I \neq 0$, there exists a minimal left ideal M such that $IM \neq 0$ (use that R is Artinian). Pick $a \in M$ such that $Ia \neq 0$; then $I(Ia) \neq 0$ and $Ia \subset M$, so $Ia = M$ by minimality of M. Thus, there exists $x \in I$ such that $a = xa$, so $1 - x$ is a zero divisor. But since $x \in J$, $1 - x$ is invertible, contradiction. □

Theorem 6.2 (Akizuki-Hopkins-Levitzki): If R is (left, right) Artinian, then R has finite length as a (left, right) module over R. In particular, R is Noetherian.

Proof. We'll show that $M_d := J^d / J^{d+1}$ is a finite length R-module. This module is annihilated by J, so it's semisimple. Recall that semisimple modules are Artinian iff they are Noetherian iff they are a finite sum of irreducibles. But J^d / J^{d+1} is Artinian, so it has a finite length. Then

$$\text{length}(R) = \sum_{i=0}^{n-1} \text{length}(M_n)$$

where the sum is finite because $J^n = 0$, so R has finite length. □

6.2 Projective covers

Definition 6.3: A module P is **projective** if $\text{Hom}(P, -)$ is exact (takes short exact sequences to short exact sequences). Equivalently, given a surjection $N \twoheadrightarrow M$, we can lift any map $P \longrightarrow M$ (non-uniquely) to a map $P \longrightarrow N$.

Example 6.4: Free modules are projective. Direct summands of projective modules are also projective, so direct summands of free modules are projective. In fact, the converse is also true, since every projective P has a surjection $R^I \twoheadrightarrow P$, so we can lift $P \cong P$ to $P \longrightarrow R^I$, which gives us a splitting of $R^I = P \oplus Q$.

Corollary 6.5: Every module is the quotient of a projective module.

Definition 6.6: A surjection $\varphi\colon M \twoheadrightarrow N$ is an **essential surjection** if for all $M' \subsetneq M$, $\varphi|_{M'}$ is not onto. That is, no proper submodule of M surjects onto N.

Definition 6.7: A **projective cover** of a module M is an essential surjection $P \twoheadrightarrow M$ from a projective module P.

Example 6.8: Let M be a finite length module and M^1 be the first term of the cosocle filtration, so $S := M/M^1 = M/JM$ is the maximal semisimple quotient (see Corollary 5.7). Then $M \twoheadrightarrow S$ is an essential surjection. One way to see this: if $N \subset M$ and $N \twoheadrightarrow S = M/JM$, then $(M/N)/J(M/N) = 0$. So by Nakayama $M/N = 0$. In fact, any essential surjection $M \twoheadrightarrow S$ with S semisimple and M finite length has this form.

Lemma 6.9:

a) Suppose $p\colon P \twoheadrightarrow M$ is a projective cover and $q\colon Q \twoheadrightarrow M$ is another surjection from a projective Q to M. Then we can write $Q \cong P \oplus Q'$ with $q|_{Q'} = 0$ and $q|_P = p$.

b) A projective cover (if it exists) is unique up to isomorphism.

Proof. b) follows from a), so it suffices to prove a). We can lift q to a map $\tilde{q}\colon Q \longrightarrow P$ with $q\colon Q \xrightarrow{\tilde{q}} P \overset{p}{\twoheadrightarrow} M$. Since p is an essential surjection, Q must be onto (as $\text{Im}(\tilde{q}) \twoheadrightarrow M$). But surjective maps between projective modules split, so we get the desired splitting of Q. □

Proposition 6.10: Suppose R is Artinian.

a) Every irreducible module has a projective cover.

b) The isomorphism classes of irreducible modules are in bijection with isomorphism classes of indecomposable projectives. This bijection sends L to its projective cover and a projective module to its cosocle (its maximal semisimple quotient).

Proof. b) follows from a): let P be an indecomposable projective. Since P is a summand of a free, there is a nonzero map from P to R, hence $P \twoheadrightarrow L$ for some irreducible L. But P_L, the projective cover of L, is a direct summand of P by Lemma 6.9, so $P \cong P_L$.

To prove a), it suffices to find a projective P_L such that $P_L/JP_L \cong L$, where $J = J(R)$, since then $P_L \twoheadrightarrow L$ is an essential surjection (see Example 6.8). We will induct on n such that $J^n = 0$. If $n = 1$, R is semi-primitive, and thus $R \cong \prod \mathrm{Mat}_{n_i}(D_i)$. Here everything is projective, so $L = P_L$. In general, we will use the lifting of idempotents; the below lemma will show that we can lift idempotents from R/I to R when $I^2 = 0$.

Suppose $n > 1$, then R/J is semi-primitive, so there exists an idempotent $\bar{e} \in R/J$ such that $(R/J)\bar{e} \cong L$. Then we can lift idempotents repeatedly along surjections $R/J^{d+1} \twoheadrightarrow R/J^d$ until we get some e in R (use Lemma 6.11 below). Then consider $P_L = Re$. This satisfies $P_L/JP_L = (R/J)\bar{e} \cong L$, and P_L is a summand of R, so we are done. □

Lemma 6.11: Let S be a ring and $I \subset S$ a 2-sided ideal such that $I^2 = 0$. Then any idempotent $e \in R := S/I$ can be lifted to an idempotent $\bar{e} \in S$.

Proof. Let e' be any lift of e, not necessarily an idempotent. We can decompose I into the direct sum

$$I = e'Ie' \oplus e'I(1-e') \oplus (1-e')Ie' \oplus (1-e')I(1-e').$$

Note that the decomposition above does not depend on the choice of e' (use that $I^2 = 0$). Notice that $\varepsilon := e'(1-e')$ lies in I (as it's 0 mod I). Moreover, it satisfies $e'\varepsilon(1-e') = (1-e')\varepsilon e' = \varepsilon^2 = 0$ (use that $I^2 = 0$), so in the direct sum decomposition ε has only nonzero first and last components. That is, we can write $\varepsilon = \varepsilon_+ + \varepsilon_-$, where $\varepsilon_+ \in e'Ie'$ and $\varepsilon_- \in (1-e')I(1-e')$. Now we claim that

$$\bar{e} := e' + \varepsilon_+ - \varepsilon_-$$

is an idempotent lifting of e. Indeed we have

$$\bar{e}(1-\bar{e}) = (e' + \varepsilon_+ - \varepsilon_-)(1 - e' - \varepsilon_+ + \varepsilon_-) = \varepsilon - e'\varepsilon_+ - \varepsilon_-(1-e') = \varepsilon - e'\varepsilon - \varepsilon(1-e') = 0.$$

□

Remark 6.12: An alternative approach to the proof of Lemma 6.11: let e' be a lift of e and set $f' := 1 - e'$. We have $1 - e'^2 - f'^2 \in I$ is nilpotent so $e'^2 + f'^2$ is invertible and it is easy to see that $e'' = \frac{e'^2}{e'^2+f'^2}$ is the desired lift of e (use that $e'^2 f'^2 = 0$).

Remark 6.13: Let P_L be the projective cover of L. Then $\mathrm{Hom}_R(P_L, L') = 0$ if $L' \neq L$, and $\mathrm{Hom}_R(P_L, L)$ is a free module over D_L^{op}, where $D_L := \mathrm{End}_R(L)$.

Corollary 6.14: Let R be an Artinian ring and write $R/J = \prod \mathrm{Mat}_{d_i}(D_i)$, $D_i = \mathrm{End}(L_i)^{\mathrm{op}}$ where the L_i are the isomorphism classes of simple R-modules and $d_i = \dim_{D_i}(L_i)$. Let P_i be the projective cover of L_i. Then

$$R \cong \bigoplus_i P_i^{d_i}$$

as a left R-module.

Proof. By Theorem 4.18, $R \cong \bigoplus_i P_i^{m_i}$ for some multiplicities m_i. Then $\mathrm{Hom}_R(R, L_i) \cong \mathrm{Hom}_R(P_i^{m_i}, L_i)$, but we have

$\mathrm{Hom}_R(R, L_i) = \mathrm{Hom}_R(R/J, L_i) = D_i^{d_i}$, while $\mathrm{Hom}_R(P_i^{m_i}, L_i) = \mathrm{Hom}_R(P_i, L_i)^{m_i} = D_i^{m_i}$, so $m_i = d_i$. □

Remark 6.15: We can view this corollary as "lifting" Theorem 2.10. Artin-Wedderburn tells us that

$$R/J \cong \bigoplus_{L_i \text{ simple}} \mathrm{End}_{D_i}(L_i)^{\mathrm{op}} \cong \bigoplus_{L_i}^{d_i},$$

thus expressing the semisimple quotient in terms of the simple modules. On the other hand, the corollary lifts this to R itself, by lifting the L_i to their projective covers:

$$R \cong \bigoplus_{L_i} P_i^{d_i}.$$

Remark 6.16: Suppose A is a finite-dimensional algebra over an algebraically closed field k. Then $\mathrm{End}_A(L) \cong k$ for all irreducible L. Then we get another proof of Theorem 3.15, as in this case, the multiplicity of L_i in M will be $\dim_k \mathrm{Hom}_A(P_L, M)$.

Corollary 6.17: Let R be an Artinian ring. Then any finitely generated R-module has a projective cover.

Proof. Induct on length. Consider $0 \to L \to M \to N \to 0$ where L is simple and suppose we know N has projective cover P_N with $\varphi\colon P_N \twoheadrightarrow N$. If $P_N \twoheadrightarrow M$, then P_N is also the projective cover of M. Otherwise, M must split as $L \oplus \mathrm{Im}(\varphi) = L \oplus N$, so $P_L \oplus P_N$ is a projective cover of M. □

6.3 Preview of Morita theory

If the P_i are the indecomposable projectives of a ring R, how is $S := \mathrm{End}_R\left(\bigoplus_i P_i^{m_i}\right)^{\mathrm{op}}$ related to R? It turns out that when $m_i \geqslant 1$, S is **Morita equivalent** to R, meaning that their module categories are equivalent.

Theorem 6.18: S is Morita equivalent to R iff $S^{\mathrm{op}} = \mathrm{End}_R(P)$, where P is a finitely generated "projective generator" of R-mod (iff $S \cong e\,\mathrm{Mat}_n(R)e$ for some n and some idempotent $e \in \mathrm{Mat}_n(R)$ satisfying $\mathrm{Mat}_n(R)e\,\mathrm{Mat}_n(R) = \mathrm{Mat}_n(R)$).

We will precisely define the projective generator next time, but when R is Artinian, it will be when $m_i \geqslant 1$ as mentioned above.

7 March 2 - Categories and Morita equivalence

Remark 7.1: We can also discuss projective covers of graded modules over graded rings. Let $R = \bigoplus_{n \geqslant 0} R_n$ with R_0 Artinian and let L be an irreducible graded module over R that is concentrated in one degree. WLOG we can assume L is concentrated in degree 0. Then $P = Re_L$ is a graded projective cover of L; $e_L \in R_0$ is the idempotent corresponding to the projective cover of L as an R_0-module.

7.1 Morita equivalence

Definition 7.2: We say that two rings are **Morita equivalent** if their categories of modules are equivalent.

(Below, we will recall some facts about categories.)

Theorem 7.3: A ring S is Morita equivalent to a ring R iff $S = \mathrm{End}_R(P)^{\mathrm{op}}$ where P is a finitely generated projective generator of the category of R-modules iff $S \cong e\,\mathrm{Mat}_n(R)e$ for some n and some idempotent $e \in \mathrm{Mat}_n(R)$ satisfying $\mathrm{Mat}_n(R)e\,\mathrm{Mat}_n(R) = \mathrm{Mat}_n(R)$.

Definition 7.4: A projective module P over a ring R is a **projective generator** if $\operatorname{Hom}(P, M) \neq 0$ for every nonzero R-module M.

Lemma 7.5: M is a generator iff R is a direct summand in M^n for some n.

Proof. If M is a generator, then for every module N, the images of all possible homomorphisms $M \to N$ generate N. This is because if S is the sum of all the images of such maps, then $\operatorname{Hom}(M, S) \to \operatorname{Hom}(M, N)$ is an isomorphism, and since M is a generator, this implies that $S \cong N$.
Now if N is finitely generated, say with generators n_i, and $n_i = \sum f_{ij}(m_j)$ where $f_{ij} \in \operatorname{Hom}(M, N)$, then only images for those finitely many f_{ij} are needed to generate N. Hence there is a surjection $M^n \twoheadrightarrow N$. In particular, if we take $N = R$, R is projective, so the surjection splits and R is a summand of M^n.
In the other direction, if R is a summand of M^n, this implies M^n is a generator, and hence M is a generator also. □

Example 7.6: R is Morita equivalent to itself. In this case, take $P = R$ (the rank 1 free module), and $R = \operatorname{End}_R(R)^{\text{op}}$. More generally, if we take $P = R^n$, then $S = \operatorname{End}_R(R^n)^{\text{op}} = \operatorname{Mat}_n(R)$ is Morita equivalent to R also.
Using the lemma, we see that if R is Artinian with indecomposable projectives $P_1, \ldots, P_n$, $P = \bigoplus P_i^{m_i}$ is a projective generator iff $m_i \geqslant 1$ for all i. In particular, if we take $m_i = 1$ for all i, then $S = \operatorname{End}_R(P)^{\text{op}}$ is what's known as a **based ring**, meaning that each irreducible L_i is a one-dimensional vector space over $D_i = \operatorname{End}_R(L_i)$.

Proposition 7.7: Let $P = R^n e$ be a finitely generated projective module, with an idempotent $e \in \operatorname{Mat}_n(R)$. Then P is a generator iff $\operatorname{Mat}_n(R) = \operatorname{Mat}_n(R) e \operatorname{Mat}_n(R)$.

Proof. Suppose $\operatorname{Mat}_n(R) = \operatorname{Mat}_n(R) e \operatorname{Mat}_n(R)$. Then we can write $1 = \sum_{i=1}^m a_i e b_i$ for $a_i, b_i \in \operatorname{Mat}_n(R)$, so the map $P^m \to \operatorname{Mat}_n(R)$ given by $(x_1, \ldots, x_m) \mapsto \sum x_i b_i$ is onto, hence we have a surjection $P^m \twoheadrightarrow \operatorname{Mat}_n(R) \twoheadrightarrow R$. So by the lemma 7.5, P is a generator.
In the other direction, suppose P is a generator. $M = \operatorname{Mat}_n(R)/\operatorname{Mat}_n(R) e \operatorname{Mat}_n(R)$ satisfies $\operatorname{Hom}_R(P, M) = \operatorname{Hom}_R(R^n e, M) = (eM)^{\oplus n} = 0$, so if $M \neq 0$, P can't be a generator. □

Note that this proves the second iff of Theorem 7.3, because for $e \in \operatorname{Mat}_n(R)$ satisfying $\operatorname{Mat}_n(R) e \operatorname{Mat}_n(R) = \operatorname{Mat}_n(R)$, we have $\operatorname{End}_R(R^n e) \cong e \operatorname{Mat}_n(R) e$.

7.2 Categories and the Yoneda Lemma

Quick review: a (small) category $\mathcal{C}$ consists of a set of objects $\operatorname{Ob}(\mathcal{C})$, a set of morphisms $\operatorname{Hom}_{\mathcal{C}}(X, Y)$ for all $X, Y \in \operatorname{Ob}(\mathcal{C})$, an identity morphism $\operatorname{id}_X \in \operatorname{Hom}(X, X)$, and an associative composition operation.

Remark 7.8: Small categories are those where $\operatorname{Ob}(\mathcal{C})$ is actually a set. Since there is no such thing as the "set of all sets", categories like Set or R-Mod are not small. We could get around this by fixing a universe and only considering sets from this universe. We could also consider "large" categories, whose objects form a collection more general than a set, called a class. We will ignore all these set-theoretic issues.

Given two categories $\mathcal{C}_1, \mathcal{C}_2$, we can talk about the category of functors $\operatorname{Fun}(\mathcal{C}_1, \mathcal{C}_2)$ whose objects are functors and whose morphisms are natural transformations.

Definition 7.9: A functor F is **faithful** if the map $\operatorname{Hom}(X, Y) \to \operatorname{Hom}(F(X), F(Y))$ is injective for all X, Y.

Definition 7.10: A functor F is **fully faithful** if the map $\operatorname{Hom}(X, Y) \to \operatorname{Hom}(F(X), F(Y))$ is an isomorphism.

Definition 7.11: A functor F is **essentially surjective** if it is surjective on isomorphism classes of objects.

Definition 7.12: A functor $F\colon \mathcal{C}_1 \to \mathcal{C}_2$ is an **equivalence of categories** if there exists $G\colon \mathcal{C}_2 \to \mathcal{C}_1$ such that $F \circ G, G \circ F$ are isomorphic to the respective identity functors (that is, they are naturally equivalent to the identity functors).

Lemma 7.13: A functor F is an **equivalence of categories** iff it is fully faithful and essentially surjective.

Proof. Since we are ignoring set-theoretic considerations, we get to use the axiom of choice. It's clear that if F is an equivalence, then it's fully faithful and essentially surjective. In the other direction, if F is essentially surjective, the axiom of choice allows us to choose $X \in \mathrm{Ob}(\mathcal{C}_2)$ and $G(X) \in \mathrm{Ob}(\mathcal{C}_1)$ such that $i_X \colon X \cong F(G(X))$. Then we can define $G(f\colon X \to Y)$ as follows: first $i_Y^{-1} \circ f \circ i_X$ gives a map $F(G(X)) \to F(G(Y))$, and because F is fully faithful, this corresponds to a unique $G(f)\colon G(X) \to G(Y)$. Then one can verify that G is indeed a functor and that $F \circ G$ and $G \circ F$ are equivalent to $\mathrm{id}_{\mathcal{C}_i}$. □

Lemma 7.14 (Yoneda Lemma): For a category $\mathcal{C}$, consider the functors $R\colon \mathcal{C}^{\mathrm{op}} \to \mathrm{Fun}(\mathcal{C}, \mathrm{Set})$ and $C\colon \mathcal{C} \to \mathrm{Fun}(\mathcal{C}^{\mathrm{op}}, \mathrm{Set})$ where $R(X)\colon T \mapsto \mathrm{Hom}(X, T)$ and $C(X)\colon T \to \mathrm{Hom}(T, X)$. Then R, C are fully faithful. Here R is for "represent" and C for "corepresent".

Proof (Sketch). For $X, Y \in \mathrm{Ob}(\mathcal{C})$, there's a natural map $\mathrm{Hom}(X, Y) \to \mathrm{Hom}(R(X), R(Y))$ given by composing with the map $X \to Y$. In the other direction, given $\varphi\colon R(X) \to R(Y)$, send it to the element $\varphi(\mathrm{id}_X) \in \mathrm{Hom}(X, Y)$. It's easy to see these are inverse bijections. The argument for C is similar. □

That is, an object in $\mathcal{C}$ is uniquely defined up to unique isomorphism up to the functor it (co)represents.

Example 7.15: The initial (resp. final) object of a category $\mathcal{C}$ is an object I (resp. F) such that $\mathrm{Hom}(I, X)$ (resp. $\mathrm{Hom}(X, F)$) is a singleton. By the Yoneda lemma, initial and final objects are unique up to unique isomorphism (if they exist). For example, in the category R-Mod, the zero module is both initial and final.

Definition 7.16: The **coproduct** (resp. **product**) is the object representing (resp. corepresenting) the product of Hom sets: $\mathrm{Hom}\left(\coprod X_i, T\right) = \prod \mathrm{Hom}(X_i, T)$ and $\mathrm{Hom}\left(T, \prod X_i\right) = \prod \mathrm{Hom}(T, X_i)$. These are unique up to unique isomorphism if they exist.

Example 7.17: In R-Mod, these both exist; coproduct is the direct sum and product is the usual product.

Remark 7.18: We can characterize the statement that a finite direct sum is the same as a finite product in categorical terms. Using the final object 0, there is a morphism $\coprod X_i \to X_i$. Hence, there is a map $\coprod X_i \to \prod X_i$, and this is an isomorphism when the X_i form a finite collection.

Remark 7.19: This can also be used to show that $\mathrm{Hom}(M, N)$ has an abelian group structure. You can define the sum of two maps $f, g\colon M \to N$ as the composition

$$M \xrightarrow{f \times g} N \times N \cong N \amalg N \xrightarrow{\mathrm{id}_N \amalg \mathrm{id}_N} N.$$

7.3 Proof of Morita equivalence theorem

Proof (of Theorem 7.3). Suppose $F\colon S\text{-Mod} \to R\text{-Mod}$ is an equivalence. We will show that $P := F(S)$ is a finitely generated projective generator in R-Mod and that $S = \mathrm{End}_S(S)^{\mathrm{op}} = \mathrm{End}_R(P)^{\mathrm{op}}$. This follows from the following observations:

- F sends projective S-modules to projective R-modules. M is projective iff $\mathrm{Hom}(M, -)$ is exact, i.e. sends a surjective map of modules to a surjective map of sets. A map of modules $T_1 \to T_2$ is surjective iff $\mathrm{Hom}(T_2, X) \hookrightarrow \mathrm{Hom}(T_1, X)$ is injective for all X. Using essential surjectivity of F, we find $N_1, N_2, Y \in S$-Mod such that $F(N_i) \cong T_i$ and $F(X) \cong Y$; then the full faithfulness of F implies that $N_1 \twoheadrightarrow N_2$. Then $\mathrm{Hom}(M, N_1) \twoheadrightarrow \mathrm{Hom}(M, N_2)$ combined with full faithfulness of F translates this into $\mathrm{Hom}(F(M), T_1) \twoheadrightarrow \mathrm{Hom}(F(M), T_2)$.
- F sends a projective generator to a projective generator, since $\mathrm{Hom}(M, N) = 0 \Leftrightarrow \mathrm{Hom}(F(M), F(N)) = 0$ by full faithfulness of F.
- F sends finitely generated projective S-modules to finitely generated projective R-modules. Use the following

characterization of finitely generated projectives: a projective P is finitely generated iff $\mathrm{Hom}(P,-)$ commutes with arbitrary coproducts (i.e. $\coprod \mathrm{Hom}(P, X_i) = \mathrm{Hom}\left(P, \coprod X_i\right)$. If P is projective and finitely generated, it's a direct summand of S^n, which has this property, so P also has this property. In the other direction, suppose $\mathrm{Hom}(P,-)$ commutes with coproducts. We know P is the direct summand of some free module, say $\bigoplus_I S$, which then splits as $P \oplus Q$. Then $\mathrm{Hom}(P, \bigoplus_I S) = \bigoplus_I \mathrm{Hom}(P, S)$, so the image of $P \hookrightarrow \bigoplus_I S$ must land in a finite direct sum $S^n = \bigoplus_J S$, $|J| < \infty$. S^n will also split as $P \oplus (Q \cap S^n)$, so P is in fact finitely generated. Since F is an equivalence of categories, it preserves the property that $\mathrm{Hom}(F(P),-)$ commutes with arbitrary coproducts, so $F(P)$ is also finitely generated projective.

Combining these three, we get that $F(S)$ is a finitely generated projective generator. Because F is fully faithful, $\mathrm{Hom}_S(S,S) \cong \mathrm{Hom}_R(F(S), F(S)) = \mathrm{End}_R(P)$, so $S = \mathrm{End}_R(P)^{\mathrm{op}}$.

In the other direction, we want to show that if $S = \mathrm{End}_R(P)^{\mathrm{op}}$ for P a finitely generated projective generator P of R-Mod, the functor $F_P \colon M \mapsto \mathrm{Hom}_R(P, M)$ is the desired equivalence of categories. Here $M \in R$-Mod and $\mathrm{Hom}_R(P, M)$ has an S-action via composition.

F_P induces an isomorphism $\mathrm{Hom}_R(P, N) \cong \mathrm{Hom}_S(F_P(P), F_P(N))$ for all N: the RHS will be $\mathrm{Hom}_S(S, F_P(N)) \cong F_P(N) = \mathrm{Hom}(P, N)$. This isomorphism coincides with the F_P-action on morphisms.

Since P is finitely generated and projective, F_P commutes with coproducts. Moreover, P is a projective generator, we claim we can find an exact sequence $P^{\oplus J} \to P^{\oplus I} \to M \to 0$.

Lemma 7.20: A projective module P is a generator iff the free module R is a direct summand in P^n for some n iff every module is a quotient of $P^{\oplus I}$.

Now we want to show that $\mathrm{Hom}(M, N) \to \mathrm{Hom}(F_P(M), F_P(N))$ is an isomorphism. Notice that if this is true for M_1, M_2, it's also true for $\mathrm{coker}(f)$, $f \colon M_1 \to M_2$ because exactness of F_P implies that both Hom-spaces are the kernel of the map $\mathrm{Hom}(M_2, N) \to \mathrm{Hom}(M_1, N)$. So by the above, it suffices to show that this is true for $M = P^{\oplus I}$, but that is what we proved above. So F_P is fully faithful.

To see that F_P is essentially surjective, take $N \in S$-Mod, which fits in an exact sequence $S^{\oplus J} \xrightarrow{f} S^{\oplus I} \to N \to 0$. Because F_P is fully faithful, $f = F_P(g)$ for $g \colon P^{\oplus J} \to P^{\oplus I}$. Hence $N \cong F_P(\mathrm{coker}(g))$. Thus, F_P is an equivalence of categories. □

Example 7.21: Now it's interesting to consider notions that are invariant under Morita equivalence. We will see that the center $Z(R)$ and cocenter $C(R)$ of a ring are such notions, i.e. if R, S are Morita equivalent, they have the same center and the same cocenter.

8 March 7 - Morita theory continued: (co)centers, functors and bimodules

8.1 Center and cocenter

Last time, we claimed that the center and cocenter are Morita invariant notions. Recall that the center $Z(R)$ is defined as

$$Z(R) := \{z \in R \mid zr = rz \ \forall r \in R\}$$

and is a commutative subring in R. The cocenter $C(R)$ is

$$C(R) := R / \sum_i [y_i, x_i],$$

i.e. the quotient of R by combinations of the commutators of elements in R. $C(R)$ is an abelian group and an $Z(R)$-module, but generally does not have a ring structure.

Proposition 8.1: If $R \sim_M S$, then $Z(R) \cong Z(S)$, $C(R) \cong C(S)$.

Remark 8.2: We will see later that $Z(R) = \mathrm{HH}^0(R)$, the 0th Hochschild cohomology, and $C(R) = \mathrm{HH}_0(R)$, the 0th Hochschild homology, and that the ith Hochschild (co)homology is also Morita invariant.

Proof. We will need several intermediate lemmas that allow us to describe $Z(R)$ and $C(R)$ purely in terms of the category of modules.

Lemma 8.3: $Z(R) \cong \mathrm{End}(\mathrm{Id}_R)$, i.e. endomorphisms of the identity functor in R-Mod, as commutative rings.

Proof. An element in $\mathrm{End}(\mathrm{Id}_R)$ is a collection of maps $z_M \in \mathrm{End}(M)$ such that $z_N \circ f = f \circ z_M$ for all $f\colon M \to N$ maps of R-modules. If we take a central element $z \in R$, it corresponds to the functor where z_M is just the left action of z on M. If we are given a collection z_M, consider z_R; note that it must commute with left multiplication by r for all $r \in R$, so it must be a central element. Hence we get the desired isomorphism. □

Definition 8.4: Let Proj_R be the category of finitely generated projective R-modules. A **trace map** for Proj_R with values in an abelian group A is an assignment of an element $\tau(P, \varphi) \in A$ for every $P \in \mathrm{Ob}(\mathrm{Proj}_R)$, $\varphi \in \mathrm{End}(P)$, such that

$$\begin{aligned} \tau(P \oplus Q, \varphi \oplus \psi) &= \tau(P, \varphi) + \tau(Q, \psi) \\ \tau(P, a \circ b) &= \tau(Q, b \circ a), a\colon Q \to P, b\colon P \to Q. \end{aligned}$$

Lemma 8.5: Let Proj_R be the category of finitely generated projective R-modules. Then $C(R)$ is the **universal abelian group receiving a trace map** for Proj_R. In other words, $C(R)$ is isomorphic (as abelian groups) to the quotient of the free abelian group generated by pairs (P, φ) by the relations $(P \oplus Q, \varphi \oplus \psi) - (P, \varphi) - (Q, \psi)$ and $(P, a \circ b) - (Q, b \circ a)$ (where $a\colon Q \to P$, $b\colon P \to Q$).

Proof. Let us restate this in terms of matrices. Let $A = (a_{ij}) \in \mathrm{Mat}_n(R)$ and set $\overline{\mathrm{Tr}}(A) = \sum a_{ii} \pmod{[R, R]}$. Then

$$\overline{\mathrm{Tr}}(AB) = \overline{\mathrm{Tr}}(BA).$$

Call the abelian group in the statement $\tilde{C}(R)$. We will use $\overline{\mathrm{Tr}}$ to construct an isomorphism $\tau\colon \tilde{C}(R) \to C(R)$. Let P be a finitely generated projective. Then it's the summand of a free, so choose Q, n such that $P \oplus Q = R^n$. Then $(\varphi \oplus 0) \in \mathrm{End}(R^n)$ with matrix A_φ. Set

$$\tau(P, \varphi) := \overline{\mathrm{Tr}}(A_\varphi).$$

Then τ is independent of choices of Q, n and satisfies $\tau(P, ab) = \tau(Q, ba)$. Also, τ is clearly additive on direct sums. So τ is a homomorphism.
It is onto since we can choose $P = R$ and φ multiplication by any element in R. To see it's injective, it suffices to show that $(R^n, A) = (R, \sum a_{ii})$ in $\tilde{C}$. But this is true because a matrix with zero sum of diagonal elements will map to 0 in $C(\mathrm{Mat}_n(R))$. □

Therefore, center and cocenter depend only on the category R-Mod, which shows they are Morita invariant. □

Example 8.6: For $a \in R$, we can consider the operator $R \to R$ of right multiplication by a. The trace of this map is just $[a] \in C(R)$.

8.2 Morita equivalence via functors and bimodules

Definition 8.7: Let R, S be rings. An R, S-bimodule M is an abelian group carrying a commuting left action of R and right action of S (i.e. a left S^{op} action). We denote such a module by ${}_R M_S$.

Given a bimodule ${}_R P_S$, we get a functor $F_P\colon S\text{-Mod} \to R\text{-Mod}$ given by $M \mapsto P \otimes_S M$. It is easy to see that $F_Q \circ F_P = F_{Q \otimes_S P}$ for bimodules ${}_R Q_S$ and ${}_S P_T$. Thus, we have a functor from R, S-Bimod $\to$ Fun $(S\text{-Mod}, R\text{-Mod})$.

Lemma 8.8: The functor $P \mapsto F_P$ is fully faithful.

Proof. There is a natural map $\mathrm{Hom}(P, Q) \to \mathrm{Hom}(F_P, F_Q)$. To construct a map in the other direction, note that $P = F_P(S)$, and this is an isomorphism of R, S-bimodules because the right action of S on $F_P(S)$ is obtained by applying F_P to $\operatorname{End}(S)$. This defines a map $\operatorname{Hom}(F_P, F_Q) \longrightarrow \operatorname{Hom}(P, Q)$, and you can check that it's the inverse bijection to the first map. □

Remark 8.9: In the proof of the Morita equivalence theorem last time, we used the functor $M \mapsto \operatorname{Hom}_R(P, M)$. This can be written as $M \mapsto \tilde{P} \otimes_R M$, where $\tilde{P} = \operatorname{Hom}_R(P, R)$ as a right R-module. We could rewrite $\operatorname{End}_R(P)^{\mathrm{op}} = \operatorname{End}_{R^{\mathrm{op}}}(\tilde{P})$. In fact, $P \mapsto \tilde{P}$ gives an equivalence of categories $\mathrm{Proj}_R^{\mathrm{op}} \longrightarrow \mathrm{Proj}_{R^{\mathrm{op}}}$.

Remark 8.10: Recall that in an equivalence of categories, you have two functors F, G and $G \circ F \simeq \mathrm{Id}_{\mathcal{C}}$, $F \circ G \simeq \mathrm{Id}_{\mathcal{D}}$. It turns out that if you fix F, G, and the first isomorphism of functors, then the second isomorphism of functors is uniquely determined so that if the two isomorphisms $F \circ G \circ F \simeq F$ coincide (from either $F \circ \mathrm{Id}_{\mathcal{C}}$ or $\mathrm{Id}_{\mathcal{D}} \circ F$, the two isomorphisms $G \circ F \circ G \simeq G$ also coincide.

Therefore, if we want to define a Morita equivalence between A, B, we can rephrase this as finding ${}_AP_B$, ${}_BQ_A$, which will give us two functors A-Mod $\longrightarrow$ B-Mod and B-Mod $\longrightarrow$ A-Mod, such that $P \otimes_B Q \simeq A$ and $Q \otimes_A P \simeq B$, i.e. their compositions are isomorphic to the respective identity functors.

Definition 8.11: A **Morita context** is the data of $A, B, {}_AP_B, {}_BQ_A$ with maps $\tau\colon P \otimes_B Q \xrightarrow{\sim} A$ and $\eta\colon Q \otimes_A P \xrightarrow{\sim} B$ such that the two arrows $P \otimes_B Q \otimes_A P \longrightarrow P$ coincide and likewise for $Q \otimes_A P \otimes_B Q \longrightarrow Q$.
This can be rewritten in matrix form: τ, η, and the bimodule structures define multiplication on matrices of the form

$$\begin{pmatrix} a & p \\ q & b \end{pmatrix}, \; a \in A,\; p \in P,\; q \in Q,\; b \in B.$$

We can now talk about pq, as $\tau(p, q)$, etc., and our compatibility condition means this matrix multiplication is associative.

Example 8.12: Let B be a ring and $M \in B$-Mod. The **derived Morita context** is given by $A = \operatorname{End}_B(M)^{\mathrm{op}}$, $Q = M$, $P = \operatorname{Hom}_B(M, B)$, and $\tau(p \otimes q) = m \mapsto p(m)q$, $\eta(q \otimes p) = p(q)$.
We can verify that the arrows $P \otimes_B Q \otimes_A P \longrightarrow P$ coincide: $p \otimes q \otimes p' \mapsto \tau(p \otimes q) \otimes p'$, which sends $m \mapsto p'(p(m)q) = p(m)p'(q)$. The other map is $p \otimes q \otimes p' \mapsto p \otimes p'(q)$, which sends $m \mapsto p(m)p'(q)$. A similar argument holds for $Q \otimes_A P \otimes_B Q \longrightarrow Q$.

Theorem 8.13: For a derived Morita context, the functors given by P, Q are inverse equivalences iff M is a finitely generated projective generator.

This is a reformulation of the theorem we proved last time. The proof is a consequence of the below lemmas.

Definition 8.14: A **generator** $M \in R$-Mod is an object such that $\operatorname{Hom}_R(M, -)$ is faithful.

Lemma 8.15: M is a generator iff for all N, there exists a surjection $M^{\oplus I} \twoheadrightarrow N$, iff R is a direct summand of M^n.

Lemma 8.16: For a derived Morita context,

a) $\tau\colon P \otimes_B Q \longrightarrow A$ is onto iff $Q = M$ is a finitely generated projective over B.

b) $\eta\colon Q \otimes_A P \longrightarrow B$ is onto iff $Q = M$ is a generator over B.

Proof. By definition $\operatorname{im}(\eta)$ is the sum of images of all homomorphisms $M \longrightarrow B$. So η is onto exactly when the sum of the images is B, which is when M is a generator. This proves b).
For a), first suppose τ is onto. Then $1_A = \mathrm{Id}_M = \sum_{i=1}^n e_i f_i$ where $f_i\colon M \longrightarrow B$ and $e_i\colon B \longrightarrow M$. Then consider the maps $m \mapsto (f_1(m), \ldots, f_n(m))$ and $(b_1, \ldots, b_n) \mapsto \sum b_i e_i$. Their composition $M \longrightarrow B^n \longrightarrow M$ is the identity, so M is a direct summand of B^n, implying it's a finitely generated projective.
In the other direction, suppose that M is a finitely generated projective. Then write $M = B^n e$ for an idempotent e. Then $\operatorname{End}(M) = e \operatorname{Mat}_n(B) e$ and we have a surjection $B^n e \otimes eB^n \longrightarrow \operatorname{End}(M)$. □

Lemma 8.17: In a Morita context, τ (resp. η) is onto implies τ (resp. η) is an isomorphism.

Proof. Suppose that $\tau \colon P \otimes_B Q \to A$ is onto. Then write $1 = \tau\left(\sum p_i \otimes q_i\right)$. Consider the map

$$Q \to Q \otimes_A P \otimes_B Q,\ q \mapsto q \otimes \left(\sum p_i \otimes q_i\right).$$

Then the composition

$$Q \to Q \otimes_A P \otimes_B Q \xrightarrow{\eta \otimes \mathrm{id}} B \otimes_B Q = Q$$

is the identity map. Tensoring with P on the left, we get the identity map $P \otimes_B Q \to P \otimes_B Q$. But the composition is also equal to

$$P \otimes_B Q \to (P \otimes_B Q) \otimes_A P \otimes_B Q \xrightarrow{\tau \otimes \mathrm{id}} A \otimes_A P \otimes_B Q = P \otimes_B Q$$

where the first arrow sends $p \otimes q \mapsto p \otimes q \otimes \left(\sum p_i \otimes q_i\right)$. Since an element in $\ker \tau$ would be killed by this composition, we must have $\ker \tau = 0$, so τ is an isomorphism. A similar argument works for η. □

8.3 Serre quotients

Motivating question: suppose that $P \in A$-Mod is a finitely generated projective but not a generator and $B = \mathrm{End}_A(P)^{\mathrm{op}}$. How are A-Mod and B-Mod related? It turns out that B-Mod is a Serre quotient of A-Mod by $\{M \mid \mathrm{Hom}(P, M) = 0\}$.

Definition 8.18: A **Serre subcategory** of an abelian category (defined next time) is a full subcategory closed under subquotients and extensions. That is, for an SES $0 \to M_1 \to M \to M_2 \to 0$, M is n the subcategory iff M_1, M_2 are.

Example 8.19: A Serre subcategory in the category of finite length modules is uniquely determined by the set of irreducible objects it contains. So such subcategories are in bijection with subsets of the set of isomorphism classes of irreducibles.

Let $\mathcal{A}$ be a Serre subcategory of an abelian category and $\mathcal{B} \subset \mathcal{A}$ a Serre subcategory.

Definition 8.20: A homomorphism $f \colon M \to N$ is an **isomorphism modulo** $\mathcal{B}$ if $\ker(f), \mathrm{coker}(f) \in \mathcal{B}$.

Definition 8.21: The **Serre quotient** $\mathcal{A}/\mathcal{B}$ is the category with a universal functor $\mathcal{A} \to \mathcal{A}/\mathcal{B}$ sending isomorphisms modulo $\mathcal{B}$ to isomorphisms. (That is, for any functor $\mathcal{A} \to C$ sending isomorphisms modulo $\mathcal{B}$ to isomorphisms, there's a unique functor $\mathcal{A}/\mathcal{B} \to C$ making the diagram commute.)

The Serre quotient has the same objects as $\mathcal{A}$, but different Hom-sets.

9 March 9 - more on Serre quotients, abelian categories

9.1 More on Serre quotients

Let $\mathcal{A}$ be a Serre subcategory in R-Mod and $\mathcal{B} \subset \mathcal{A}$ be a Serre subcategory. We defined the Serre quotient abstractly, but here is a more concrete description:

- Objects of $\mathcal{A}/\mathcal{B}$ are objects of $\mathcal{A}$.
- The morphisms $\mathrm{Hom}_{\mathcal{A}/\mathcal{B}}(M, N)$ are equivalence classes of "roof diagrams" $M \leftarrow M' \to N$, where the left arrow $M \leftarrow M'$ is an isomorphism modulo $\mathcal{B}$ (i.e. its kernel and cokernel are both in $\mathcal{B}$). Two roof diagrams $M \leftarrow M' \to N$ and $M \leftarrow M'' \to N$ are equivalent if there exists a map $M'' \to M'$ commuting with the other arrows, i.e.

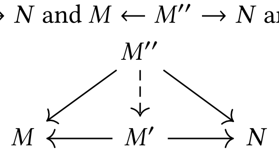

Another way to phrase this:

$$\mathrm{Hom}_{\mathcal{A}/\mathcal{B}}(M,N) = \operatorname*{colim}_{M'\to M} \mathrm{Hom}(M',N)$$

where the colimit is taken over the category of objects $M' \in \mathcal{A}$ equipped with isomorphisms modulo $\mathcal{B}$ to M.

Remark 9.1: We could also phrase $\mathrm{Hom}_{\mathcal{A}/\mathcal{B}}$ in terms of "lower roof" diagrams, where the arrows are reversed, so $\mathrm{Hom}_{\mathcal{A}/\mathcal{B}}(M,N)$ consists of diagrams $M \to N' \leftarrow N$ where $N' \leftarrow N$ is an isomorphism modulo $\mathcal{B}$. Why are these definitions equivalent? Given a lower roof diagram, you can construct the upper roof by setting $M' := M \times_{N'} N$ (the pullback), i.e. $\ker(M \oplus N \to N')$. Given an upper roof diagram, you can set N' to be the pushforward, namely $N' := \mathrm{coker}(M' \to M \oplus N)$.

Example 9.2: Let $\mathcal{A}$ be the category of finite length modules over an Artinian algebra R and $\mathcal{B}$ be the subcategory of modules that do not have some fixed irreducibles $L_1, \ldots, L_i$ in their Jordan-Holder series. Then $\mathcal{A}/\mathcal{B}$ will be the category of finite length modules over

$$S = \mathrm{End}\left(\bigoplus_{j=i+1}^{n} P_j\right)^{\mathrm{op}},$$

i.e. the sum of the projective covers of the remaining irreducibles.
If we remove the finite length assumption, then you get a special case of

$$R\text{-Mod}/(P^{\perp}) \cong \mathrm{End}_R(P)^{\mathrm{op}}\text{-Mod}.$$

Example 9.3: Let R be a commutative ring, $\mathcal{A} = R$-Mod, $I \subset R$, and $\mathcal{B}$ be the modules where every element of I acts locally nilpotently. Then $\mathcal{A}/\mathcal{B} = \mathrm{QCoh}(\mathrm{Spec}(R)\backslash Z_I)$ where Z_I is the zero set of I. This is a quasiaffine scheme (not necessarily affine).
You can also get (quasi)coherent sheaves on more general varieties using the Serre quotient. For example, a projective variety X over a field k can be obtained as $\mathrm{Proj}(A)$ for a positively graded commutative algebra A with $A_0 = k$. Then $\mathrm{Coh}(X)$ is the Serre quotient

$$A\text{-Mod}^{\mathrm{gr}}_{\mathrm{fg}}/A\text{-Mod}_0$$

where $A\text{-Mod}^{\mathrm{gr}}_{\mathrm{fg}}$ is the category of finitely generated graded modules and $A\text{-Mod}_0$ is the subcategory of finite-dimensional (equivalently, concentrated in finitely many degrees) modules. Geometrically, this corresponds to starting with dilation equivariant sheaves on the cone $\mathrm{Spec}(A)$ and throwing away the origin.

9.2 Adjoint functors and (co)limits

Definition 9.4: An **adjunction** for a pair of functors $L\colon C_1 \to C_2$, $R\colon C_2 \to C_1$ is an isomorphism

$$\mathrm{Hom}_{C_2}(L(X), Y) \cong \mathrm{Hom}_{C_1}(X, R(Y))$$

that is functorial in X, Y. Then we say that L, R are **adjoint functors**, that L is the **left adjoint** of R, and R is the **right adjoint** of L.
The Yoneda Lemma indicates that L determines R up to unique isomorphism and vice versa (if it exists).

Example 9.5: In general, free and forgetful functors are adjoint; for example, the functor sending a set S to the corresponding free structure (group, abelian group, module, algebra, etc.) on S is left adjoint to the forgetful functor to Set. Likewise, the functor sending a Lie algebra to its universal enveloping algebra $\mathfrak{g} \to U(\mathfrak{g})$ is left adjoint to the functor sending an associative algebra to itself but as a Lie algebra (with the Lie bracket $[x,y] = xy - yx$).

Example 9.6: It is possible for a functor to have a left adjoint but no right adjoint: for example, the full embedding of commutative rings into associative rings has a left adjoint sending R to the quotient by the 2-sided ideal generated by the commutators. But it has no right adjoint.

Example 9.7 (Tensor-Hom adjunction): Let ${}_AP_B$ be a bimodule. Then $L = P \otimes_B -: B\text{-Mod} \to A\text{-Mod}$ is left adjoint to $R = \mathrm{Hom}_A(P, -): A\text{-Mod} \to B\text{-Mod}$.

Example 9.8: Consider the category $\mathrm{Fun}(\mathcal{D}, \mathcal{C})$ of functors from $\mathcal{D} \to \mathcal{C}$. The functor

$$\mathrm{Cons}: \mathcal{C} \to \mathrm{Fun}(\mathcal{D}, \mathcal{C}),\ \mathrm{Cons}(X)(Y) = X, \mathrm{Cons}(f) = \mathrm{Id}_X$$

has right adjoint Cons^*; this may or may not exist, but if it does, $\mathrm{Cons}^*(F)$ is the **limit** or **inverse limit** of F. Likewise, the left adjoint ${}^*\mathrm{Cons}$, if it exists, sends F to the **colimit** or **direct limit** of F.
We can describe the limit more concretely. $\mathrm{Cons}^*(F)$ has the following property: it is the universal object equipped with compatible maps into $F(i), i \in \mathcal{D}$. That is, given an object X with compatible maps into $F(i)$ for $i \in \mathcal{D}$, there is a unique map $X \to \mathrm{Cons}^*(F)$ that makes the diagram commute.
Note that not all limits and colimits may exist in a category. Some examples of limits and colimits: the colimit of $\bullet \quad \bullet$ is the coproduct, while its limit is the product. The limit of $\bullet \to \bullet \leftarrow \bullet$ is the pullback, the colimit of $\bullet \leftarrow \bullet \to \bullet$ is the pushout.

9.3 Additive categories

We are interested in categories like R-Mod that have additional structure.

Definition 9.9: An **additive category** $\mathcal{A}$ is a category where each Hom set has the structure of an abelian group such that the composition is bilinear and the following properties hold:

a) There exists an object $0_{\mathcal{A}}$ such that $\mathrm{Hom}(0_{\mathcal{A}}, 0_{\mathcal{A}}) = 0$ (the zero group),

b) For every $M_1, M_2 \in \mathcal{A}$, there exists an object $S = M_1 \oplus M_2$ with morphisms $p_i: S \to M_i$ and $\iota_i: M_i \to S$ such that $p_1\iota_2 = p_2\iota_1 = 0$, $p_1\iota_1 = \mathrm{id}_{M_1}, p_2\iota_2 = \mathrm{id}_{M_2}$, and $\iota_1 p_1 + \iota_2 p_2 = \mathrm{id}_S$.

This implies that $\mathrm{Hom}(0, M) = \mathrm{Hom}(M, 0) = 0$, so 0 is both the initial and final object. Also, S is both the coproduct and product of M_1, M_2: you can see this by noting that the corresponding fact is true for abelian groups, then apply this to $\mathrm{Hom}(S, X)$ and $\mathrm{Hom}(X, S)$.

Notice that we were able to deduce a global property (about Hom in every object) from a local property (only looking at $M, N, S, 0$).

Note 9.10: We don't need to include an addition on Hom sets in the definition. If we know that there is an initial and final object and that therefore, the resulting map from coproducts to products is an isomorphism, you can recover addition on Hom sets, as discussed in the category of modules. But it's more convenient to list it in the definition.

9.4 Abelian categories

An abelian category is essentially a "category where you can do homological algebra" and was introduced by Grothendieck.

Definition 9.11: An **abelian category** is an additive category satisfying

- AB1: existence of kernel and cokernels: that is, objects representing the functor $X \to \ker(\mathrm{Hom}(X, M) \to \mathrm{Hom}(X, N))$ and corepresenting the functor $X \to \ker(\mathrm{Hom}(N, X) \to \mathrm{Hom}(M, X))$ for a morphism $f: M \to N$. Morphisms with zero kernel are **monic** and morphisms with zero cokernel are **epic**.
- AB2: A monic morphism is a kernel; that is, for $f: M \to N$, let $K = \ker f$ and $C = \operatorname{coker} f$, then $\operatorname{coker}(K \to M) \to \ker(N \to C)$ is an isomorphism.

One can also add the additional axioms

- AB3: the existence of arbitrary coproducts
- AB4: the coproduct of any family of monic morphisms is monic

A subobject of A is an object A_i with a monic morphism $A_i \hookrightarrow A$. The sum of some subobjects A_i is $\mathrm{im}\left(\coprod A_i \to A\right)$. The intersection of two subobjects A, B of C is $\ker(C \to C/B \oplus C/A)$. We can add one last axiom

- AB5: $(\sum A_i) \cap B = \sum (A_i \cap B)$ for a collection of increasing subobjects A_i in A.

We can also define AB3,4,5*: a category satisfies ABn* if $\mathcal{A}^{\mathrm{op}}$ satisfies ABn.
If a category satisfies AB1-5, it's called a **Grothendieck category**.

Definition 9.12: A category $\mathcal{D}$ is **filtered** if $\mathrm{Ob}(\mathcal{D}) \neq \varnothing$ and for all $a, b \in \mathcal{D}$, there exists $c \in \mathcal{D}$ such that $\mathrm{Hom}(a, c), \mathrm{Hom}(b, c)$ are nonempty and such that for every pair of parallel morphisms $e, f: a \to b$, there exists $g: b \to c$ such that $ge = gf$.

Remark 9.13: An equivalent definition of a filtered category is that a category is filtered if and only if colimits over the category commute with finite limits (into the category of sets). Therefore, filtered colimits generalize the properties we expect from (generalized) intersections and unions.

The key feature of Grothendieck categories is that filtered colimits exist and are exact.

Remark 9.14: The category of R-modules satisfies AB5, AB3*, and AB4*.

Remark 9.15: The only abelian category satisfying AB3-5 and AB3*-5* is the zero category. Sketch of proof: consider an object X in such a category and let Σ, Π be the coproduct and product of countably many copies of X. There is a canonical map $c: \Sigma \to \Pi$; it is monic because it's the colimit of embeddings of a direct summand and epic since it is the inverse limit of surjections to a direct summand. Hence c is an isomorphism. Now consider the composition φ of the arrows $X \to \Pi \xrightarrow{c^{-1}} \Sigma \to X$ where the first arrow is the diagonal and the second arrow is the codiagonal. Then one can check that $\varphi + \mathrm{id}_X = \mathrm{id}_X$ because "$\infty + 1 = \infty$". Hence $\mathrm{id}_X = 0$ and $X \cong 0$.

9.5 Compact projective generators and Serre quotients revisited

Definition 9.16: An object M is **compact** if $\mathrm{Hom}(M, -)$ commutes with filtered colimits.

If M is projective, this follows from commuting with arbitrary direct sums, since

$$\mathrm{colim}(F) = \operatorname{coker}\left(\bigoplus_{e:\, a \to b} F(a) \to \bigoplus_{a} F(a)\right)$$

where the morphism takes $x \mapsto x - F(e)(x)$. In general, this is not true, though it is true that every compact module is finitely generated.

Definition 9.17: An object P is a **generator** if $T \mapsto \mathrm{Hom}(P, T)$ is a faithful functor. For a projective object, this is equivalent to the Definition 7.4. Alternatively, we could say that if $P^\perp$ is the full subcategory whose objects are M such that $\mathrm{Hom}(P, M) = 0$, then a projective object P is a generator iff $P^\perp \cong \{0\}$.

Theorem 9.18: An abelian category with coproducts (satisfying AB3) and a projective compact generator is $\mathrm{End}(P)^{\mathrm{op}}$-Mod where P is a projective compact generator.

The proof is the same as the proof in the Morita theory case.

Corollary 9.19: Let P be a compact projective object in an AB3 abelian category $\mathcal{A}$. Let $\mathcal{B} = P^\perp$. Then $\mathcal{A}/\mathcal{B} \cong \mathrm{End}(P)^{\mathrm{op}}$-Mod.

Proof (Sketch). It's clear that 1) P is projective in $\mathcal{A}/\mathcal{B}$ (use the lower roof diagram Homs) and 2) P is a generator (in $\mathcal{A}/\mathcal{B}$). $\mathcal{B}$ is closed under coproducts, so the projection functor $\mathcal{A} \to \mathcal{A}/\mathcal{B}$ commutes with coproducts. Hence P is compact in $\mathcal{A}/\mathcal{B}$. □

This proves the claim at the beginning of Section 8.3.

References for this lecture include the original article [9], which still makes for excellent reading. Textbook expositions can be found in [13] and in the appendix to [18].

10 March 14 - Exts and Tors, Resolutions

10.1 Ext and Tor

Definition 10.1: Let M, N be objects in an abelian category. $\mathrm{Ext}^i(M, N)$ is the derived functor of Hom. Recall that Hom is left exact in the second argument and right exact in the first argument, so you can take either the right derived functor of $\mathrm{Hom}(M, -)$ or the left derived functor of $\mathrm{Hom}(-, N)$, and these are the same. Although the most useful formalism for this is the derived category, we can also work in a typical category.

The key property of Ext is that it is the **universal delta functor**. Delta functors were introduced by Grothendieck in his Tohoku paper; essentially, they turn short exact sequences to long exact sequences. Given a short exact sequence $0 \to L \to M \to N \to 0$, a delta functor is a functorial family of (additive) functors T^i with boundary morphisms $\delta_i \colon T^i(N) \to T^{i+1}(L)$ such that $0 \to T^0(L) \to T^0(M) \to T^0(N) \xrightarrow{\delta_0} T^1(L) \to \cdots \to$ is exact. One can define morphisms of delta functors as families of natural transformations that commute with the boundary morphisms, and a delta functor is universal when giving a morphism to any other delta functor is equivalent to only giving the natural transformation in degree zero.

To show that something is a universal functor, it's enough to show that it's "effaceable" (in the language of Grothendieck), meaning that every element $\varphi \in \mathrm{Ext}^i(M, N), i > 0$ is killed by some injection $N \hookrightarrow N'$.

To actually compute Ext, we use projective and injective resolutions. A projective resolution of M is an exact sequence $\cdots \to P_2 \to P_1 \to M \to 0$ where the P_i are projective; these always exist for R-modules. $\mathrm{Ext}(M, N)$ is computed by applying $\mathrm{Hom}(-, N)$ to the resolution, removing M, and computing the cohomology of the resulting complex. You can also compute Ext using left injective resolutions of N, i.e. $0 \to N \to I_1 \to \cdots \to$ where I_i are injective.

Then, in this case, it's easy to see that Ext is effaceable – every N has an injection into an injective I, and every M receives a surjection from a projective P, so these maps efface all elements in $\mathrm{Ext}^i, i > 0$ because $\mathrm{Ext}^i(P, N) = \mathrm{Ext}^i(M, I) = 0$ for $i > 0$.

A better formal setting for this is the homotopy category of complexes $\mathcal{Ho}(R)$. The morphisms in this category are defined as follows: for C_1, C_2 complexes in R-Mod, let $\mathrm{Hom}^\bullet(C_1, C_2)$ be the complex where $\mathrm{Hom}^i(C_1, C_2) = \prod_j \mathrm{Hom}(C_1^i, C_2^{i+j})$ and define $\mathrm{Hom}_{\mathcal{Ho}(R)}(C_1, C_2) := H^0(\mathrm{Hom}^\bullet(C_1, C_2))$. $\mathrm{Hom}^\bullet$ has a differential, which is to take the supercommutator with d. That is, it consists of maps $f \colon C_1 \to C_2$ that commute with d modulo the equivalence that $f \sim g$ if $f - g = d_{C_2} h + h d_{C_1}$ where $h \colon C_1^i \to C_2^{i+1}$ is any collection of maps.

Exercise 10.2: There is a full embedding R-Mod $\to \mathcal{H}o(R)$ taking $M \mapsto P_M$, a projective resolution of M, which is unique up to unique isomorphism in $\mathcal{H}o(R)$. (That is, projective resolutions are "unique up to homotopy").

Let $\mathcal{H}o^0(R)$ be category of complexes of projectives in nonpositive degree with $H^i = 0, i < 0$ (so they are exact outside of degree 0). Then there is an equivalence $\mathcal{H}o^0(R) \to R$-Mod taking $C \mapsto H^0(C)$.

Remark 10.3: M is projective iff $\mathrm{Ext}^1(M, N) = 0$ for all N. If M is projective, it has projective resolution $0 \to M \to M \to 0$. If $\mathrm{Ext}^1(M, N) = 0$, then $0 \to N \to P \twoheadrightarrow M \to 0$ has a splitting for all N (either use the definition that Ext^1 is in bijections with extensions or use the long exact sequence), so M is projective.

10.2 Projective, injective, and homological dimension

Definition 10.4: The **projective dimension** $\mathrm{pdim}(M)$ of a module M is

$$\mathrm{pdim}(M) := \max\left\{i \mid \exists N \text{ s.t. } \mathrm{Ext}^i(M, N) \neq 0\right\} \in \mathbb{Z}_{\geqslant 0} \cup \{\infty\}.$$

If $0 \to M' \to P \twoheadrightarrow M \to 0$, $\mathrm{pdim}(M) = \mathrm{pdim}(M') + 1$ unless M is projective. This is because for $i \geqslant 1$, the LES says $0 \to \mathrm{Ext}^i(M', N) \to \mathrm{Ext}^{i+1}(M, N) \to 0$.
Alternately, we can define projective dimension as the length of the minimal projective resolution. For example, if $\mathrm{pdim}(M) = 1$, that means $0 \to Q \to P \to M \to 0$ is a resolution of M.

Definition 10.5: The **injective dimension** $\mathrm{idim}(M)$ of a module N is

$$\mathrm{idim}(M) := \max\left\{d \mid \exists M \text{ s.t. } \mathrm{Ext}^d(M, N) \neq 0\right\}$$

or the length of the minimal injective resolution.

Definition 10.6: The **homological dimension** $\mathrm{hdim}(R)$ of a ring R is the maximal projective dimension of an R-module, which is the same as the maximal injective dimension of an R-module. It is also

$$\mathrm{hdim}(R) := \max\left\{d \mid \exists M, N \text{ s.t. } \mathrm{Ext}^d(M, N) \neq 0\right\}.$$

Remark 10.7: N is injective iff $\mathrm{Ext}^1(M, N) = 0$ for all cyclic M.

Proof. We'll show that if $0 \to M' \hookrightarrow M$ and $M' \to N$, we can extend this to a map $M \to N$. By Zorn's Lemma, it suffices to show that it can be extended to some $M'' \subset M$ with $M' \subsetneq M''$. Pick $m \in M \setminus M'$ and let M'' be the submodule generated by M', m. We have an exact sequence $0 \to M \to M'' \to M''/M' \to 0$ and by construction M''/M' is cyclic.
Hence, if $\mathrm{Ext}^1(M''/M', N) = 0$, there exists an extension of $M' \to N$ to $M'' \to N$. □

This remark implies that

$$\mathrm{hdim}(R) = \max\left\{d \mid \mathrm{Ext}^d(M, N) \neq 0 \text{ for some } M, N \text{ s.t. } M \text{ is f.g.}\right\}.$$

Example 10.8: If R is left Noetherian, a finitely generated M has a resolution of finitely generated projectives. Then $\mathrm{Ext}^i(M, -)$ commutes with filtered colimits. Hence, we can assume that N is also finitely generated in the above definition of homological dimension. If R is Artinian, we can say more: it suffices to consider only irreducible M, N.

10.3 Cartan matrices

In this section suppose that R is Artinian and R-Mod refers only to finitely generated modules. If R has finite homological dimension, then K^0 (R-Mod) (Definition 3.17) is generated by classes of projective modules: for every simple, write a projective resolution $0 \to P_L^i \to \cdots \to P_L^0 \to L \to 0$, then $[L] = \sum (-1)^i P_L^i$.

Definition 10.9: Let $L_1, \ldots, L_n$ be the irreducibles for a ring R, and $P_1, \ldots, P_n$ be their projective covers. The **Cartan matrix** of R is the $n \times n$ matrix with $C_{ij} = [P_j : L_i]$, the multiplicity of L_i in P_j.
If R is finite-dimensional over an algebraically closed field, we can also say that $C_{ij} = \dim_k \operatorname{Hom}(P_i, P_j)$.

We then get an identification $K^0(R\text{-Mod}) \cong \mathbb{Z}^n$ via $[M] \mapsto ([M : L_1], \ldots, [M : L_n])$. Hence, if R has finite homological dimension, $C \in \mathrm{GL}_n(\mathbb{Z})$; the ijth entry of C^{-1} is

$$\sum_d (-1)^d \# \left\{ \text{summands of } P_{L_i}^d \text{ isomorphic to } P_j \right\}.$$

Corollary 10.10: If $n = 1$, R has finite homological dimension iff $R = \mathrm{Mat}_m(D)$ for a skew field D and some integer m.

Now let R be Artinian and M a finitely generated module. We will characterize **minimal projective resolutions**.

Lemma 10.11: Let $\cdots \xrightarrow{d_{-2}} P^{-1} \xrightarrow{d_{-1}} P^0 \to M \to 0$ be a projective resolution and set $C^i = \ker(d_i) = \operatorname{im}(d_{i-1})$. Then TFAE:

a) $P^{-i-1} \xrightarrow{d_{-i-1}} C^{-i}$ is a projective cover for every i.
b) $L \otimes_R P^\bullet$ has 0 differential for all irreducible right R-modules L.
c) $\operatorname{Hom}_R(P^\bullet, L)$ has 0 differential for all irreducible L.

Resolutions satisfying these properties are minimal. From *a*), if it exists, it is unique up to non-unique isomorphism because projective covers are unique. From *c*), we see that in the minimal resolution,

$$P^{-d} = \bigoplus_i P_i^{m_i^d}, \ \operatorname{Ext}^d(M, L_i) = D_i^{m_i^d}$$

where $m_i^d = \dim_{D_i}(\operatorname{Ext}^d(M, L_i))$.

Proof. $P \twoheadrightarrow M$ is a projective cover iff it induces an isomorphism $\operatorname{Hom}(M, L) \to \operatorname{Hom}(P, L)$ for all irreducibles L.
First, $P \twoheadrightarrow M$ iff $\operatorname{Hom}(M, L) \hookrightarrow \operatorname{Hom}(P, L)$ for all irreducibles L. If $P \twoheadrightarrow M$, then by applying $\operatorname{Hom}(-, L)$, which is left exact, we see that $\operatorname{Hom}(M, L) \hookrightarrow \operatorname{Hom}(P, L)$. If the map $P \to M$ is not onto, then $\operatorname{coker}(P \to M)$ is nonzero finitely generated, so it has irreducible quotient L. Then $M \twoheadrightarrow L$ is in the kernel of $\operatorname{Hom}(M, L) \to \operatorname{Hom}(P, L)$, so this map is not injective.
If $P \twoheadrightarrow M$ is a projective cover, and there exists $P \to L$ that doesn't come from some $M \to L$, then $\ker(P \to L) \twoheadrightarrow M$, so the surjection is not essential, a contradiction. If $P \twoheadrightarrow M$ is not a projective cover, then there exists $Q \hookrightarrow P$ with $Q \twoheadrightarrow M$. Then P/Q has a simple quotient L, and the map $P \twoheadrightarrow P/Q \twoheadrightarrow L$ cannot come from a map $M \to L$: if it did, then $P \twoheadrightarrow M \to L$ should pull back to $Q \twoheadrightarrow M \to L$, but this is the zero map because it's also the composition $Q \hookrightarrow P \twoheadrightarrow P/Q \twoheadrightarrow L$, which is zero. Hence $\operatorname{Hom}(M, L) \to \operatorname{Hom}(P, L)$ is not surjective.
By definition $0 \to P^{-i-1} \to P^{-i} \to C^{-i+1} \to 0$. If $P^{-i} \to C^{-i+1}$ is a projective cover, then $\operatorname{Hom}(C^{-i+1}, L) \cong \operatorname{Hom}(P^{-i}, L)$, iff $\operatorname{Hom}(P^{-i}, L) \cong \operatorname{Hom}(P^{-i+1}, L)$. □

Remark 10.12: This generalizes to $\mathbb{Z}_{\geqslant 0}$-graded rings where A_0 is Artinian and A_d is finitely generated over A_0. A common setting where this appears is an algebra A over an algebraically closed field k where A_0 is semisimple and A_d is finite-dimensional over k. In this setting, there are still indecomposable projectives. In minimal resolutions, each term has finitely many generators in each degree. The graded irreducibles are concentrated in one degree (use that if M is a graded A-module, then $M_{\geqslant k} := \bigoplus_{i \geqslant k} M_i \subset M$ is a A-submodule of M for any $k \in \mathbb{Z}$). It follows that graded irreducible A-modules are annihilated by $A_{\geqslant 1}$ so they are just irreducible A_0-modules (up to a shift of grading).
If $A_0 = k$ is just a field, and for finitely generated (graded) M, we can consider its Poincare series $\sum_i \dim(M_i) t^i \in \mathbb{Z}((t))$. More generally, if A_0 is semisimple then one can consider series $P_M := \sum_i [M_i] t^i \in \mathbb{Z}((t))^n$ where n is the number of irreducibles for A_0 and $[M_i] \in K^0$ (A_0-Mod) $\cong \mathbb{Z}^n$. The Cartan matrix C now lies in $\mathrm{GL}_n(\mathbb{Z}[[t]])$ instead of $\mathrm{Mat}_n(\mathbb{Z})$ (it is clear that $C \in \mathrm{Mat}_n(\mathbb{Z}[[t]])$ and $C(0) = \mathrm{Id}$ since A_0 is semisimple, it then follows that $C \in \mathrm{GL}_n(\mathbb{Z}[[t]])$). If L_i has finite homological dimension, and A is Noetherian then $C^{-1} \in \mathrm{Mat}_n(\mathbb{Z}[t])$.
For example, if $A = k[x]$, considered as a graded algebra with $\deg x = 1$, then $n = 1$, $L_1 = k$, $P_1 = k[x]$, so $C = \sum_{i=0}^{\infty} t^i = \frac{1}{1-t}$, and $C^{-1} = 1 - t$.

11.1 More on the Hattori-Stallings Dennis trace

Recall from Lemma 8.5 that the cocenter $R/[R,R] = C(R)$ receives a universal trace map $\tau(P, \varphi) \in C(R)$ where P is a finitely generated projective and $\varphi \in \mathrm{End}(P)$. In fact, if R is Noetherian and of finite homological dimension, you can extend τ to $\tau(M, \varphi)$ where M is any finitely generated module. To do so, choose a finite projective resolution $0 \to P^{-n} \to P^{-n+1} \to \cdots \to P^0 \to M \to 0$ (which exists because R has finite homological dimension). Then we can lift φ to $\tilde{\varphi} \in \mathrm{End}(P_M^\bullet)$ and this will be unique up to homotopy. Define

$$\tau(M, \varphi) = \sum_i (-1)^i \tau(P^{-i}, \tilde{\varphi}^{-i})$$

which is well-defined because $M \mapsto P_M^\bullet$ is a fully faithful functor to the homotopy category of complexes. Moreover, τ is additive on short exact sequences of modules.

Corollary 11.1: If R is a finite-dimensional algebra of finite homological dimension over an algebraically closed field k, then $J(R) \subset [R,R]$.

Proof.

Lemma 11.2: For $M \in R$-Mod and $\varphi \in \mathrm{End}_R(M)$, we can find a φ-invariant Jordan-Holder series of M.

Proof. Consider $\varphi|_{\mathrm{Soc}(M)} : \mathrm{Soc}(M) \to \mathrm{Soc}(M)$, where $\mathrm{Soc}(M) = \bigoplus_i L_i^{d_i}$ is the socle of M. Then φ induces an R-linear map $L_i^{d_i} \to L_i^{d_i}$ i.e. an element of $\mathrm{End}_R(L_i^{d_i}) = \mathrm{Mat}_{d_i}(k)$ (use Schur's lemma) and this matrix has an eigenvector, which generates a φ-invariant irreducible submodule in M. Then by inducting on the length of M, we get a φ-invariant Jordan-Holder series. □

Thus, $\tau(M, \varphi) = \sum_i \tau(L_i, \lambda_i) = \sum_i \lambda_i \tau(L_i, 1)$ where $\lambda_i \in k$. It follows that the elements $\tau(L_i, 1) \in C(R)$ generate $C(R)$ as a vector space over k (use Lemma 8.5 or Example 8.6). We conclude that $C(R)$ has dimension (over k) at most the number of irreducibles L_i. On the other hand, let $\bar{R} := R/J(R)$ and note that $C(R) \twoheadrightarrow C(\bar{R})$. It's easy to see that $C(\bar{R}) = k^{\# L_i}$, so $C(R) \cong C(\bar{R})$ and $J(R) \subset [R,R]$. □

Question 11.3: Is there a way to prove this without using the trace map?

11.2 Minimal resolutions and Koszul rings

Given a module M, how can we find its minimal resolution? For certain algebras called Koszul algebras, their minimal resolutions are called Koszul complexes. One great reference is [5, Section 2].

Let A be a nonnegatively graded algebra over an algebraically closed field k with A_0 semisimple. We will be interested in the case $A_0 = k$ so we can write $A = k \oplus A_{>0}$.

Remark 11.4: An elementary property of minimal resolutions for graded modules is that if $M = \bigoplus_{i \geqslant 0} M_i$ has a minimal resolution $P^{-n} \to \cdots \to P^{-1} \to P^0 \to M \to 0$ by graded projectives, then P^{-i} must be concentrated in degrees i and higher, since the projective cover $P \twoheadrightarrow M$ is an isomorphism in the bottom degree. Since A_0 is semisimple and M_0 is an A_0-module, we have $M_0 = \bigoplus_i L_i^{\oplus n_i}$ is a direct sum of simple A_0-modules. Now since for a simple module (of a semisimple ring) the projective cover is just the simple module itself, and $P_0^0 \to M_0$ is the first step in a minimal projective resolution, by Lemma 10.11 we find that $P_0^0 = M_0$, and hence all other $P_0^{-i} = 0$. Continuing in this fashion, we see that at each step, the lowest degree of P^{-i} maps isomorphically onto the lowest available module in the kernel of the previous map (as the lowest available module, it is an A_0-module, hence repeating the previous argument) which is necessarily in degree $\geq i$, by induction.

We will need the following technical lemma.

Lemma 11.5: Let M be a finitely generated graded module over A. Then the following properties are equivalent:
(i) M is generated by degree i elements,
(ii) $M \otimes_A k$ is concentrated in degree i,
(iii) $\mathrm{Hom}_A(M, k)$ is concentrated in degree $-i$.

Proof. Lemma follows from the Nakayama lemma together with the fact that

$$\mathrm{Hom}_A(M, k) = \mathrm{Hom}_{A_0}(M/A_{>0}M, k) = (M/A_{>0}M)^*.$$

□

Definition 11.6: We say that A is **Koszul** if P^{-i} is generated by degree i elements. Equivalently, $\mathrm{Tor}_i^A(k, k)$ (where each of the k are in degree 0) is concentrated in degree i, which is equivalent to $\mathrm{Ext}_A^i(k, k)$ is concentrated in degree $-i$ (use Lemma 11.2 above).

Theorem 11.7:

a) Koszul rings are **quadratic**, i.e. $A = T(V)/\langle I \rangle$, where $T(V)$ is the tensor algebra for a vector space V and I is a subspace of $V \otimes V$.

b) If A is Koszul, then $\mathrm{Ext}_A^\bullet(k, k) = A^!$, where $A^!$ is the **dual quadratic algebra** $T(V^*)/\langle I^\perp \rangle$.

Remark 11.8: One major reason why Koszul rings are important is that if A is Koszul, then we have a derived equivalence

$$D^b(A\text{–mod}) \simeq D^b(A^!\text{–mod}).$$

Example 11.9: Let $A = T(V)$, so $I = 0$. Then the dual quadratic algebra is $A^! = T(V^*)/\langle V^* \otimes V^* \rangle = k \oplus V^*$. Hence $\mathrm{Ext}_A(k, k)$ is only nonzero in degrees 0 and 1. $k = T(V)/\langle V \rangle$ then has a free resolution in degrees 0 and 1.

Example 11.10: Let $A = \mathrm{Sym}(V) = T(V)/\langle \wedge^2 V \rangle$. Then $A^! = T(V^*)/\langle \mathrm{Sym}^2(V^*) \rangle = \bigwedge^\bullet V^*$.

Definition 11.11: The d**th Veronese subalgebra** $A^{(d)}$ is $\bigoplus_{n=0}^\infty A_{nd}$.

Let us mention the following theorem without a proof (see [3] for details).

Theorem 11.12: If A is a finitely generated commutative algebra, $A^{(d)}$ is Koszul for large d.

Remark 11.13: Using the approach of [6, Section 2] or [16] (see also Remark 12.2 below) one can easily prove (using Serre's vanishing theorem) that for every $m \in \mathbb{Z}_{\geqslant 0}$ and large enough d (depending on m) the algebra $A^{(d)}$ has the following property: P^{-i} is generated by degree i elements for $i \leqslant m$. The statement of Theorem 11.12 is stronger, and the proof is more involved.

11.3 Koszul complexes

Remark 11.14: Assume $A = T(V)/\langle I\rangle$ is quadratic. Then

$$A_n = T^n(V)/\langle I\rangle_n = V^{\otimes n}/\left(\sum_{i=0}^{n-2} V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2}\right).$$

Define

$$R_n := \bigcap_{i=0}^{n-2} V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2}$$

to be the intersection rather than the sum. Then $R_n = (A^!_n)^*$,

$$R_n^* = V^{*\otimes n}/\left(\sum_{i=0}^{n-2} (V^*)^{\otimes i} \otimes I^\perp \otimes (V^*)^{\otimes n-i-2}\right) = A^!_n. \tag{1}$$

Definition 11.15: The **Koszul complex**, denoted $\mathbb{K}^\bullet$, is a complex of free A-modules $\cdots \to A \otimes_k R_2 \to A \otimes_k R_1 \to A$. As (graded) vector spaces, $\mathbb{K}^\bullet = \bigoplus_{n=0}^\infty \mathbb{K}^\bullet_n$. The differential of $\mathbb{K}^\bullet_n$ is given by:

$$\mathbb{K}_n^{i-n} = A_i \otimes R_{n-i} \hookrightarrow A_i \otimes V \otimes R_{n-i-1} \longrightarrow A_{i+1} \otimes R_{n-i-1} = \mathbb{K}_n^{i+1-n}$$

where the left map is induced by the natural embedding $R_{n-i} \subset V \otimes R_{n-i-1}$ and the right map is induced by the multiplication $A_i \otimes V \longrightarrow A_{i+1}$.

Definition 11.16: Let V be a vector space. A **distributive lattice** of subspaces of V is a collection of subspaces satisfying

- For Y in the lattice, $X \subset Y$ is also in the lattice
- For X, Y in the lattice, $X + Y$ is also in the lattice
- For X, Y, Z in the lattice, $X \cap (Y + Z) = (X \cap Y) + (X \cap Z)$ (distributivity).

Theorem 11.17 (Theorem 11.7 cont.):

a) Koszul rings are **quadratic**, i.e. $A = T(V)/\langle I\rangle$, where $T(V)$ is the tensor algebra for a vector space V and I is a subspace of $V \otimes V$.

b) If A is Koszul, then $\mathrm{Ext}^\bullet_A(k,k) = A^!$, where $A^!$ is the **dual quadratic algebra** $T(V^*)/\langle I^\perp\rangle$.

c) Say A is a quadratic algebra. It is Koszul iff $\mathbb{K}$ is exact, i.e. $H^i(\mathbb{K}) = 0$ for all $i \neq 0$, iff $\mathbb{K}$ is the minimal resolution of the left module k.

d) Say A is a quadratic algebra. It is Koszul iff for all n, the $n-1$ vector spaces $V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2}$, $i = 0, \ldots, n-2$, generate a distributive lattice of subspaces of $V^{\otimes n}$.

Lemma 11.18: A collection of vector subspaces in a vector space W generate a distributive lattice iff there exists a basis of W such that every subspace is spanned by a subset of the basis.

Proof. Clear. □

Remark 11.19: The distributive property for the subspaces of $V^{\otimes n}$ described above is what implies the exactness of $\mathbb{K}_n$. Moreover, the exactness of $\mathbb{K}_m$, $m \leqslant n$, implies the distributive property for the subspaces of $V^{\otimes n}$.

For a collection $\mathcal{W} = (W; W_1, \ldots, W_n)$, where W is a vector space and $W_1, \ldots, W_n \subset W$ are its subspaces let $K^{-l} = K^{-l}(\mathcal{W}) := \bigcap_{i=1}^{l-1} W_i / \left((W_{l+1} + \ldots + W_n) \cap \left(\bigcap_{i=1}^{l-1} W_i \right) \right)$, where $l = 0, 1, \ldots, n+1$.

For example, we have

$$K^{-n-1} = \bigcap_{i=1}^{n} W_i,\ K^{-n} = \bigcap_{i=1}^{n-1} W_i,\ K^{-n+1} = \bigcap_{i=1}^{n-2} W_i / \left(W_n \cap \left(\bigcap_{i=3}^{n} W_i \right) \right), \ldots, K^{-1} = W / \sum_{i=2}^{n} W_i, K^0 = W / \sum_{i=1}^{n} W_i.$$

We have the natural maps $K^l \to K^{l+1}$ that make $K^\bullet = K^\bullet(\mathcal{W})$ into a complex.

Lemma 11.20: If $W_1, \ldots, W_n \subset W$ are proper subspaces and every proper subset of $\{W_1, \ldots, W_n\}$ generate a distributive lattice then $W_1, \ldots, W_n$ do the same iff $K^\bullet(\mathcal{W})$ is exact.

Proof. It is clear that if $\{W_1, \ldots, W_n\}$ generate a distributive lattice then $K^\bullet(\mathcal{W})$ is exact (for example, use Lemma 11.18).
Assume now that $K^\bullet(\mathcal{W})$ is exact. We prove the claim by the induction on n. We follow [4, Section 4.5].
We will use the following notations. Given a collection $U_1, \ldots, U_n \subset U$, say that a subspace $B \subset U$ is a splitting for $(U; U_1, \ldots, U_n)$ if there exists $C \subset U$ such that $B \oplus C = U$ and $(B \cap U_i) + (C \cap U_i) = U_i$. We will say that $(U; U_1, \ldots, U_n)$ is indecomposable if U has no proper nonzero subspaces that split $(U; U_1, \ldots, U_n)$. The following easy facts will be extremely useful.

Fact (1): The subspace $U_1 \cap \ldots \cap U_i$ or $U_1 + \ldots + U_i$ is a splitting for $(U; U_1, \ldots, U_n)$ iff it is a splitting for $(U; U_{i+1}, \ldots, U_n)$.

Proof. Clear. □

Fact (2): Assume that $(U_1 + \ldots + U_i) \cap (U_{i+1} \cap \ldots \cap U_j) = 0$ and $U_{i+1} \cap \ldots \cap U_j$ is a splitting for $(U; U_1 + \ldots + U_i, U_{j+1}, \ldots, U_n)$. Then $U_{i+1} \cap \ldots \cap U_j$ is a splitting for $(U; U_1, \ldots, U_n)$.

Proof. Let $(U_{i+1} \cap \ldots \cap U_j) \oplus B$ be a splitting for $(U; U_1 + \ldots + U_i, U_{j+1}, \ldots, U_n)$. Our goal is to check that it also gives a splitting for $(U; U_1, \ldots, U_n)$. From $(U_1 + \ldots + U_i) \cap (U_{i+1} \cap \ldots \cap U_j) = 0$ we conclude that $U_1 + \ldots + U_i \subset B$ so $U_1, \ldots, U_i \subset B$. It remains to check that $U_k = (U_k \cap (U_{i+1} \cap \ldots \cap U_j)) + (U_k \cap B)$ for $k = i+1, \ldots, j$. This is clear since $U_{i+1} \cap \ldots \cap U_j \subset U_k$. □

Fact (2'): Assume that $(U_1 \cap \ldots \cap U_i) \cap (U_{i+1} + \ldots + U_j) = 0$ and $U_1 \cap \ldots \cap U_i$ is a splitting for $(U; U_{i+1} + \ldots + U_j, U_{j+1}, \ldots, U_n)$. Then $U_1 \cap \ldots \cap U_i$ is a splitting for $(U; U_1, \ldots, U_n)$.

Proof. Same proof as the one of Fact 2. □

Let us now return to the proof. Without losing the generality, we can assume that $\mathcal{W} = (W; W_1, \ldots, W_n)$ is indecomposable and all W_i are nonzero (and proper).
It then follows (use that by the inductive assumption, $W_1 \cap W_2, W_3, \ldots, W_n \subset W$, $W_1, \ldots, W_{n-2}, W_{n-1} + W_n \subset W$ form distributive lattices and then apply Fact 1) that:

$$W_1 \cap W_2 = 0,\ W_{n-1} + W_n = W. \tag{2}$$

We can assume that $n \geqslant 4$ (for $n = 3$ the statement is clear, use exactness of $K^\bullet(\mathcal{W})$).
Assume that $n = 4$. We have $W_1 \cap W_3 \cap W_4 = 0 = W_2 \cap W_3 \cap W_4$ (use Fact 1). We also have

$$(W_1 + W_2) \cap W_3 \cap W_4 = ((W_1 + W_2) \cap W_3) \cap ((W_1 + W_2) \cap W_4) = ((W_1 \cap W_3) + (W_2 \cap W_3)) \cap ((W_1 \cap W_4) + (W_2 \cap W_4)).$$

We claim that the intersection $((W_1 \cap W_3) + (W_2 \cap W_3)) \cap ((W_1 \cap W_4) + (W_2 \cap W_4))$ is zero. Indeed, if $a + b = c + d$ for some $a \in W_1 \cap W_3$, $b \in W_2 \cap W_3$, $c \in W_1 \cap W_4$, $d \in W_2 \cap W_4$ then $a - c = d - b$ must lie in $W_1 \cap W_2 = 0$ i.e. $a = c \in W_1 \cap W_3 \cap W_4 = 0$, $d = b \in W_2 \cap W_3 \cap W_4 = 0$ so $a = b = c = d = 0$. We conclude that $(W_1 + W_2) \cap W_3 \cap W_4 = 0$. It then follows from Fact 2 that $W_3 \cap W_4$ splits $(W; W_1, W_2, W_3, W_4)$ so we must have $W_3 \cap W_4 = 0$ i.e. $W = W_3 \oplus W_4$. It remains to note that $W = W_3 \oplus W_4$ is splitting for $(W; W_1, W_2, W_3, W_4)$, and a contradiction finishes the argument.

If $n > 4$. The propperty (2) implies that $(W; W_1, \ldots, W_n)$ remains acyclic after arbitrary transpositions of $W_1, \ldots, W_{n-2}$ (by acyclic, we mean that the corresponding complex $K^\bullet$ is exact, it will be equal to zero in this case). So we may assume that for certain $1 \leqslant i \leqslant n-3$ one has $A = W_1 \cap \ldots \cap W_i \neq 0$ and each $i+1$-tuple from $W_1, \ldots, W_{n-2}$ intersects by zero. Put $B = U_{i+1} + \ldots + U_{n-2}$. Then $(W; A; B; W_{n-1}, W_n)$ satisfies the assumptions of Lemma 11.20 (acyclicity follows from the fact that $A \cap B = 0$ and $W_{n-1} + W_n = W$) so (from $n = 4$ case) we conclude that $A; B; W_{n-1}, W_n \subset W$ generate a distributive lattice so A is a splitting for $(W; W_1, \ldots, W_n)$ by Fact $2'$. Since $A \neq 0$, we get a contradiction.

□

Proof (of Theorem 11.17). If $\mathrm{Tor}_1(k,k)$ is concentrated in degree 1, then $A_{\geqslant 1}$ is generated by degree 1 elements as an A-module (use the exact sequence $0 \to A_{\geqslant 1} \to A \to k \to 0$ together with Nakayama). Hence, A is generated by degree 1 elements as a ring. Let $V = A_1$ and write $A = T(V)/I$. We have a map $A \otimes V \to A$. Using that $\mathrm{Tor}_2(k,k)$ is concentrated in degree 2, we see that $\ker(A \otimes V \to A)$ is generated by elements in $A_1 \otimes V = V \otimes V$. These elements considered as elements of $V \otimes V \subset T(V)$ generated the ideal $\ker(T(V) \twoheadrightarrow A)$, so A is quadratic.

Exactness of Koszul complex implies Koszul: If $\mathbb{K}_n$ is exact for $n \geqslant 1$, then $\mathbb{K}$ is a free resolution of k as an A-module. So now we can use it to compute $\mathrm{Ext}^\bullet_A(k,k)$. Since $R^*_n \xrightarrow{0} R^*_{n-1}$ and $R^*_n = A^!_n$, $\mathrm{Ext}^n_A(k,k) = A^!_n$. You also have to check that this is compatible with multiplication, but after showing that, we can deduce that A is Koszul. To be continued next lecture. For the compatibility with multiplication, see §13.2. □

12 March 21 - Koszul rings cont., bar complex

12.1 Finishing up Koszul rings

Proof (of Theorem 11.17, cont.) Subspaces $V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2} \subset V^{\otimes n}$, $i = 0, 1, \ldots, n-2$ generate a distributive lattice iff $\mathbb{K}^\bullet_n$ is exact: to see that it is enough to note that $\mathbb{K}^\bullet_n = K^\bullet(\mathcal{W})$ for

$$\mathcal{W} = (V^{\otimes n}; V^{\otimes n-2} \otimes I, V^{\otimes n-3} \otimes I \otimes V, \ldots, V^{\otimes n-i-1} \otimes I \otimes V^{\otimes i-1}, \ldots, I \otimes V^{\otimes n-2}).$$

Now the claim follows from Lemma 11.20 (using induction on n).
It's easy to see that $\mathbb{K}_n$ acyclic implies that $\mathbb{K}$ is a resolution for the trivial module, and $\mathrm{Tor}^A_i(k,k)$ is concentrated in degree i, so A is Koszul. In the other direction, suppose A is Koszul. We will inductively check acyclicity in the first d terms of the complex, which looks like $\cdots \to A \otimes I \to A \otimes V \to A$. If this complex is exact up to degree d, then the minimal space of generators for $\ker(A \otimes R_d \to A \otimes R_{d-1})$ is (some lift of) $\mathrm{Tor}^A_{d+1}(k,k)$. Because A is Koszul, this is in degree $d+1$, so it's a subspace in $A_1 \otimes R_d = V \otimes R_d$. It is the kernel of the multiplication map, so it must be R_{d+1}, so we're done. □

Remark 12.1: In commutative algebra, a "Koszul complex" often refers to a complex formed given a commutative ring R and n elements $x_1, \ldots, x_n \in R$. The last arrow in the complex is $R^{\oplus n} \to R$, sending $r_1, \ldots, r_n \mapsto \sum_{i=1}^n x_i r_i$. The Koszul complex for $\mathrm{Sym}(V)$ is an example of this.

Remark 12.2: We are now ready to give a sketch of the proof of the fact that for every $m \in \mathbb{Z}_{\geqslant 0}$, and large enough d, the algebra $A^{(d)}$ has the following property: P^{-i} is generated by degree i elements for $i \leqslant m$ (see Remark 11.13 above). So, our goal is to check that for every $n \in \mathbb{Z}_{\geqslant 0}$ the degree nth term of the Koszul complex for $A^{(d)}$ is exact for large enough d.
First of all, we can assume that A is generated by $A_1 = V$. Set $X := \mathrm{Proj}\, A$. We can assume that the natural morphism $X \hookrightarrow \mathbb{P}^N$ is a closed embedding. We have a natural (very ample) line bundle $\mathcal{O}_X(1)$ on X with $\Gamma(X, \mathcal{O}_X(1)) = A_1 = V$. Set $Y := X^n$, $\mathcal{L} := \mathcal{O}_X(1)^{\boxtimes n}$. For a closed $Z \subset Y$ we have $H^0(Y, \mathcal{L}) = V^{\otimes n}$ and denote by $Q_Z \subset V^{\otimes n}$ the kernel of $H^0(Y, \mathcal{L}) \to H^0(Z, \mathcal{L})$. Let $\Delta_i \subset X^n$ be the diagonal given by $x_i = x_{i+1}$. We have $Q_{\Delta_{i+1}} = V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2}$.
Let S^n be the (finite) set of closed subschemes of Y generated by $\{\Delta_i \mid i = 1, \ldots, n-1\}$ and $X^n, \varnothing$ via unions and (scheme-theoretic) intersections. Using Serre's vanishing theorem, we can assume that the statements of [6, Corollary 1.7] are satisfied for S^n. It then follows from [6, Lemma 2.1] that subspaces $V^{\otimes i} \otimes I \otimes V^{\otimes n-i-2}$ generate a distributive lattice of subspaces of $V^{\otimes n}$ so we are done by (the proof of) Theorem 11.17 (d).

Corollary 12.3: The **Poincare series** of a graded algebra is

$$P_A(t) = \sum d_n t^n, d_n = \dim A_n.$$

If A is Koszul, then $P_A(t)P_{A^!}(-t) = 1$.

Proof. This follows from the (graded) Euler characteristic of $\mathbb{K}$. If you look degree by degree, you can find that the Euler characteristic of $\mathbb{K}_n$ is the nth coefficient of $P_A(t)P_{A^!}(-t)$ (see (1)) so the total Euler characteristic of $\mathbb{K}$ is equal to $P_A(t)P_{A^!}(-t)$. Recall now that the Euler characteristic of $\mathbb{K}_n$ can also be computed as the alternating sum of dimensions of the cohomology of $\mathbb{K}_n$. It remains to note that $\mathbb{K}_n$ is exact for $n > 0$ and $\mathbb{K}_0 = k$ (sitting in degree 0). It follows that the total graded Euler characteristic of $\mathbb{K}$ is equal to 1. □

Example 12.4: Let $A = \bigwedge^n V$. Then $P_A(t) = (1+t)^n$. Likewise, $P_{\mathrm{Sym}(V)} = \frac{1}{(1-t)^n}$.

Proof (of Theorem 11.17, cont. again). Finally, we need to check that $A^! \simeq \mathrm{Ext}^\bullet_A(k,k)$ is an algebra isomorphism. First, we explain how to make $\mathrm{Ext}^\bullet$ into an algebra: $\mathrm{Ext}^\bullet_A(k,k) = H^*(\underline{\mathrm{Hom}}(P^\bullet, P^\bullet))$ for a projective resolution $P^\bullet$; $\underline{\mathrm{Hom}}$ is a DGA (differential graded algebra).
Here is how $A^!$ acts on $\mathbb{K}$: start with the action of $T(V^*)$ on $T(V)$ by contracting tensors $V^{*\otimes i} \times V^{\otimes n} \longrightarrow V^{\otimes n-i}$. Restrict this to $V^{*\otimes i} \times R_n \longrightarrow R_{n-i}$, which factors through $A^!_i \times R_n$. Recall that $\mathbb{K}^{-n} = A \otimes R_n$. Consider the map

$$(A \otimes R_n) \otimes A^!_i \longrightarrow A \otimes R_{n-i} = \mathbb{K}^{-(n-i)}.$$

This is the $A^!$-action, and it commutes with the differential. Moreover, for $a \in A^!$, the composition $\mathbb{K} \xrightarrow{a} \mathbb{K} \longrightarrow k$ represents the class of a. Hence, this is an algebra isomorphism. □

Remark 12.5: Let Proj_A be the projective graded A-modules. Then $A^!$ gives us an equivalence of derived categories

$$\mathcal{H}o(\mathrm{Proj}_A^{f.g.}) \simeq \mathcal{H}o(\mathrm{Proj}_{A^!}^{f.g.})$$

sending $M(1) \mapsto M[1](-1)$ where $M(1)_i = M_{i+1}$ and $[\cdot]$ is some homological stuff we won't discuss here. The idea is to use k as a generator for the derived category and consider the functor $F_k\colon M \longrightarrow \mathrm{RHom}(k, M)$ which generalizes $F_P(M) = \mathrm{Hom}(P, M)$.

Remark 12.6: Let A_1, A_2 quadratic, $A_i = T(V_i)/I_i$. Then

$$A_1 \otimes A_2 = T(V_1 \oplus V_2)/I_1 \oplus I_2 \oplus \langle v_1 \otimes v_2 - v_2 \otimes v_1 \rangle$$

and

$$(A_1 \otimes A_2)^! = T(V_1^* \oplus V_2^*)/I_1^\perp \oplus I_2^\perp + \langle v_1 \otimes v_2 + v_2 \otimes v_1 \rangle,$$

the "super" (signed) tensor product.

Remark 12.7: If A is commutative and $I \supset \wedge^2(V)$, then $I^\perp \subset S^2V^*$. Then all the relations of $A^!$ will be relations between anticommutators and $A^!$ will be the enveloping algebra of a Lie superalgebra.

For more on Koszul rings, see [4] and [5].

12.2 Bar complex and Hochschild (co)homology

Definition 12.8: Let A be any algebra over a field k. Then the **bar complex** of A, denoted by $\beta(A, A)$, is

$$\cdots \longrightarrow A \otimes_k A \otimes_k A \longrightarrow A \otimes_k A \longrightarrow 0,$$

so that $\beta(A, A)_n = A^{\otimes_k (n+2)}$, where the maps are

$$d\colon a_0 \otimes \cdots \otimes a_n \mapsto a_0 a_1 \otimes a_2 \otimes \cdots \otimes a_n - a_0 \otimes a_1 a_2 \otimes \cdots \otimes a_n + \cdots + a_0 \otimes a_1 \otimes \cdots \otimes a_{n-1} a_n.$$

The RHS is also written as $a_0|a_1|\cdots|a_n$. Then $d^2 = 0$.

Lemma 12.9: The bar complex $\beta(A, A)$ is an exact resolution of A for any associative algebra A (where the last map $A \otimes_k A \longrightarrow A$ is given by $a \otimes b \mapsto ab$).

Proof. The map $h\colon a_0 \otimes \cdots \otimes a_n \mapsto 1 \otimes a_0 \otimes \cdots \otimes a_n$ satisfies $dh + hd = \mathrm{id}$, so it is a chain homotopy. (These are co-simplicial maps.) □

The bar complex is also a complex of A-bimodules. The left action is on a_0, and the right action is on a_n. A is the regular A-bimodule (i.e., $A \otimes_k A^{\mathrm{op}}$-module), and all the other terms are free, so the bar complex is a free resolution for A by $A \otimes_k A^{\mathrm{op}}$-modules. This allows us to compute $\mathrm{Ext}^i_{A \otimes A^{\mathrm{op}}}(A, A)$ and $\mathrm{Tor}_i^{A \otimes A^{\mathrm{op}}}(A, A)$.

The bar complex also gives us a free resolution of every A-module by tensoring with M. The cohomology of the bar complex is $\mathrm{Tor}_i^A(A, M) = 0$ for $i > 0$.

Definition 12.10: The **Hochschild homology** of A is the homology of the bar resolution of A by $A \otimes_k A^{\mathrm{op}}$-modules, i.e. $\mathrm{HH}_i(A) := \mathrm{Tor}_i^{A \otimes_k A^{\mathrm{op}}}(A, A) = H_i(\beta(A, A) \otimes_{A \otimes_k A^{\mathrm{op}}} A)$. The **Hochschild cohomology** $\mathrm{HH}^i(A)$ of A is defined to be the cohomology of $\mathrm{Hom}_{A \otimes_k A^{\mathrm{op}}}(\mathrm{Bar}, A)$, so the nth term is $A \otimes (A^{\otimes n})^*$; this is equal to $\mathrm{Ext}^i_{A \otimes_k A^{\mathrm{op}}}(A, A)$. If A is graded, you can likewise define graded Hochschild cohomology.

Remark 12.11: The identifications with Tor and Ext are true when A is a flat k-module, for example when k is a field. In general, Hochschild homology and cohomology and equal to *relative* Tor and Ext, respectively: $\mathrm{HH}_\bullet(A) = \mathrm{Tor}_\bullet^{A \otimes_k A^{\mathrm{op}}/k}(A, A)$ and $\mathrm{HH}^\bullet(A) = \mathrm{Ext}^\bullet_{A \otimes_k A^{\mathrm{op}}/k}(A, A)$.

Remark 12.12: If A is augmented, you can use the reduced bar complex; let A_+ be the augmentation ideal, the reduced bar complex has terms $A \otimes_k A_+ \otimes_k \ldots \otimes_k A_+ \otimes_k A$. This allows you to compute $\mathrm{Ext}^i_A(k, k)$ and $\mathrm{Tor}_i^A(k, k)$, and indeed $A^!$ is in the bottom degree. Furthermore, more generally, we can talk about Hochschild homology of A-modules, sheaves, and even categories.

Remark 12.13: More concretely, we can describe Hochschild homology and the cohomology of the complex

$$C_n = A^{\otimes_k (n+1)}, \quad d : a_0 \otimes \cdots \otimes a_n \mapsto a_0 a_1 \otimes \cdots \otimes a_n - a_0 \otimes a_1 a_2 \otimes \cdot \otimes a_n + \cdots \pm a_n a_0 \otimes a_1 \otimes \cdots \otimes a_{n-1}.$$

This is most obvious when A is commutative, as then $A \cong A^{\mathrm{op}}$ and $A \otimes_{A_e} \beta(A,A)_n = A \otimes_{A \otimes_k A} A^{\otimes_k (n+2)} = A^{\otimes_k (n+1)} = C_n$.
Note that this complex has an extra "wrap-around" term compared to the bar complex, and also ends with $C_0 = A$. Let's see where this extra term comes from. We have an identification

$$\varphi : C_n(A) \xrightarrow{\sim} \beta(A,A)_n \otimes_{A \otimes_k A^{\mathrm{op}}} A$$

given as follows. First, note that this last term A in the RHS acts on $\beta(A,A)_n$ by A-action on the first term of $\beta(A,A)_n$ and A^{op}-action on the last term of $\beta(A,A)_n$. The map is given by

$$\varphi : a_0 \otimes \cdots a_n \mapsto [a_0 \otimes \cdots \otimes a_n \otimes 1 \otimes_{A_e} 1].$$

Now when we apply d on the term $[a_0 \otimes \cdots \otimes a_n \otimes 1 \otimes 1]$, we get

$$\begin{aligned} d(\varphi(a)) = d([a_0 \otimes \cdots \otimes a_n \otimes 1 \otimes 1]) =& [a_0 a_1 \otimes a_2 \otimes \cdots \otimes a_n \otimes 1 \otimes 1] \\ & - [a_0 \otimes a_1 a_2 \otimes \cdots \otimes a_n \otimes 1 \otimes 1] \\ & \cdots \\ & + (-1)^{n-1} [a_0 \otimes \cdots \otimes a_{n-2} \otimes a_{n-1} a_n \otimes 1 \otimes 1] \\ & + (-1)^n [a_0 \otimes a_1 \otimes \cdots \otimes a_{n-1} \otimes a_n \otimes 1]. \end{aligned}$$

Note that these terms, except for the last one, are precisely φ of the differential in the bar complex, i.e. $\pm\varphi(a_0 \otimes \cdots \otimes a_i a_{i+1} \otimes \cdots \otimes a_n)$. However, the last term gives us:

$$\begin{aligned} [a_0 \otimes \cdots \otimes a_n \otimes_{A_e} 1] &= [a_0 \otimes \cdots \otimes a_{n-1} \otimes 1 \otimes_{A_e} a_n], \\ &= [a_n a_0 \otimes a_1 \otimes \cdots \otimes a_{n-1} \otimes 1 \otimes 1], \\ &= \varphi(a_n a_0 \otimes a_1 \otimes \cdots \otimes a_{n-1}), \end{aligned}$$

using the $A \otimes_k A^{\mathrm{op}}$-action.
Similarly, Hochschild cohomology can be described as the cohomology of the complex $C^n = \mathrm{Hom}(A^{\otimes_k n}, A)$.

13.1 Deformations and Hochschild cohomology

From the definition of Hochschild (co)homology, we see that $\mathrm{HH}_0 = C(A) = A/[A,A]$ the cocenter and $\mathrm{HH}^0 = \mathrm{Hom}_{A \otimes A^{\mathrm{op}}}(A,A) = Z(A)$ the center.

We also have a nice description for HH^1: the kernel of d is $\{\varphi \colon A \to A \mid \varphi(ab) = \varphi(a)b + a\varphi(b)\}$ and the image of d is $\{\varphi \mid \exists x \text{ s.t. } \varphi(a) = [a,x]\}$. So HH^1 is the derivations modulo the inner derivations, i.e., the outer derivations of A.

Lemma 13.1: $\mathrm{HH}^2(A)$ is in bijection with isomorphism classes of first order deformations of A.

Definition 13.2: An n**th order deformation** of A is an algebra $\tilde{A}$ free over $k[t]/(t^{n+1})$ and an isomorphism $\tilde{A}/t\tilde{A} = A$. A **formal deformation** of A is the same as above, but over $k[[t]]$ (and we need to use flatness instead of free), and a **polynomial deformation** of A is the one over $k[t]$.

Proof. Suppose $\tilde{A}$ is a first order deformation of A and fix an isomorphism $\tilde{A} \simeq A \otimes_k (k[t]/(t^2))$. The multiplication μ on $\tilde{A}$ will correspond to a cocycle: it is determined by $\mu(a,b)$ for $a, b \in A$, and we must have $\mu(a,b) = ab$ modulo

t, so we can say that $\mu(a,b) = ab + \varphi(a,b)t$ where $\varphi\colon A\otimes A \to A$. Then associativity of μ corresponds to φ being a cochain since we need

$$a\varphi(b,c) - \varphi(ab,c) + \varphi(a,bc) - \varphi(a,b)c = 0.$$

Given any cocycle, we can define a deformation of A by defining multiplication on $A\otimes k[t]/t^2$ to be $ab + \varphi(a,b)t$. An isomorphism of deformations $\widetilde{A_\varphi} \simeq_f \widetilde{A_\psi}$ is a map $f\colon \tilde{a} \mapsto \tilde{a} + tf(a)$ for $f\colon A \to A$, since again it only depends on the values it takes on A. So f is an algebra homomorphism iff

$$(\psi - \varphi)(a,b) = af(b) - f(a)b,$$

that is, if $\psi - \varphi$ is a coboundary. □

Remark 13.3: Given an nth order deformation, the obstruction to extending it to an $(n+1)$st order deformation lies in $\mathrm{HH}^3(A)$; an expression in terms of the multiplication on $\tilde{A}$ must vanish in HH^3. Hence, if $\mathrm{HH}^3(A) = 0$, any deformation can be extended, and the set of all such extensions is in bijection with HH^2. However, this bijection is not canonical. Exercise: to get a canonical bijection, you also need the data of a torsor over HH^2.

Example 13.4: What is $\mathrm{HH}_\bullet(A)$ and $\mathrm{HH}^\bullet(A)$ for $A = k[x_1,\dots,x_n] = \mathrm{Sym}(V)$? For simplicity, assume char $k = 0$. We see we need to compute

$$\mathrm{HH}^\bullet(A) = \mathrm{Ext}^\bullet_{\mathrm{Sym}(V\oplus V)}(\mathrm{Sym}(V), \mathrm{Sym}(V))$$

and we already know how to do this: change coordinates using the Koszul complex to find that it's $\mathrm{Sym}(V)\otimes \bigwedge(V^*)$.
In particular, we remarked above that HH^1 is the outer derivations. For a commutative ring, there are no inner derivations, so $\mathrm{HH}^1(A)$ is exactly the derivations of $\mathrm{Sym}(V)$, which are

$$\left\{\sum_{i=1}^n p_i\partial_{x_i}\right\}, p_i \in k[x_1,\dots,x_n], \partial_{x_i}\colon P \to \frac{\partial P}{\partial x_i}.$$

Hence, $\mathrm{HH}^\bullet(A)$ is the polyvector fields on $V^* = \mathrm{Spec}(\mathrm{Sym}(V))$ and $\mathrm{HH}_\bullet(A) \simeq \mathrm{Sym}(V)\otimes\bigwedge V$, $\bigwedge V$ is in degree -1. These are the differential forms on V, Ω^i is in degree $-i$.

Remark 13.5: Hochschild-Kostant-Rosenberg generalized this to a smooth algebraic variety V. $\mathrm{HH}_\bullet$ and $\mathrm{HH}^\bullet$ carry more structure, related to differential geometry: the de Rham differential on forms corresponds to the Connes differential, which corresponds to cyclic cohomology. The latter uses the fact that the differential in the bar complex has cyclic symmetry.
The polyvector fields have a Schouten bracket, extending the commutator of vector fields $[v,w] = \mathrm{Lie}_v(w)$ (the Lie derivative). This generalizes to $\mathrm{HH}^\bullet(A)$, e.g. the obstruction in HH^3 for extending the 1st order deformation is $[h,h]$ where h is the deformation class.

13.2 Cobar complex and $A^!$

Let A be an augmented algebra, $A = k\cdot 1 \oplus A_+$, $A_+ = \bigoplus_{n\geqslant 1} A_n$. This induces a splitting of the bar resolution

$$\beta(A,A)_n = \left(A\otimes \underbrace{A_+\otimes\cdots\otimes A_+}_{n}\otimes A\right)\oplus\bigoplus \mathrm{span}(a_0\otimes \underbrace{1\otimes 1\otimes\cdots\otimes 1}_{\in A_+^{\otimes n}}\otimes a_n),$$

where we denote

$$\overline{\beta}(A,A)_n = A\otimes\underbrace{A_+\otimes\cdots\otimes A_+}_{n}\otimes A$$

to be the **reduced bar complex**. This is because $d(\alpha\otimes 1\otimes\beta) = d(\alpha)\otimes 1\otimes\beta \pm \alpha\otimes 1\otimes d(\beta)$ + stuff and you can check that the stuff is all like $\cdots a_{i-1}\otimes a_i\cdots - \cdots a_{i-1}\otimes a_i\cdots$ so it cancels. Therefore, both of the above are closed under the d-action. Hence, we can consider the reduced bar resolution $\overline{\beta}(A, A)$ and we can use it to compute graded $\operatorname{Ext}^\bullet_A(k, k)$ and show that it is $A^!$.

Define the graded dual of $M = \bigoplus M_i$ to be $M^* := \bigoplus M_i^*$; in this notation, the cobar complex is

$$A_+^* \to A_+^{*\otimes 2} \to \cdots$$

where the first is in degree $\leqslant -1$, the second is in degree $\leqslant -2$, and so on. Consider the degree $-i$ part in the ith term; it will equal $(V^*)^{\otimes i}$ where $V = A_1$, and

$$\operatorname{Ext}^i_A(k, k)_{-i} \simeq V^*/d()$$

where $d()$ is spanned by $d(a_1 \otimes a_2 \otimes \cdots \otimes b \otimes a_j \otimes \cdots \otimes a_i)$ where $a_k \in V^*$ and $b \in A_2^*$; this is

$$\pm d(a_1 \otimes a_2 \otimes \cdots \otimes db \otimes \cdots \otimes a_i).$$

So $d\colon A_2^* \to A_1^* \otimes A_1^*$, $A_2 = A_1 \otimes A_1/I$, and I is the space of degree 2 relations. $A_2^* = I^\perp \overset{d}{\hookrightarrow} V^* \otimes V^*$. So $V^*/d() \simeq A_i^!$, the quadratic dual to the quadratic part of A.

The cobar complex above is a DGA acting on the bar resolution of k. Hence, $A^! \simeq \bigoplus \operatorname{Ext}^i_A(k, k)_{-i}$ is an algebra isomorphism. This gives us another way to see that $\operatorname{Ext}^\bullet_A(k, k) \cong A^!$, but the advantage of this approach is that we can see that this is an isomorphism not just of graded vector spaces, but of graded algebras as well.

Note 13.6: For our next topic, we'll need that $H^*(G, M) = \operatorname{Ext}^\bullet_{\mathbb{Z}[G]}(\mathbb{Z}, M)$ where G is a group (see §15.2 below).

13.3 Central simple algebras and Brauer group

We will now turn our attention to an important class of algebras called central simple algebras.

Definition 13.7: We say that R is a **central simple algebra** (csa) over a field k if R is simple, Artinian, and its center is k.

Note that such rings are of the form $R = \operatorname{Mat}_n(D)$ for D a skew field. The center of D is a field k. We want to understand central simple algebras of finite dimension over a given field k.

Theorem 13.8:
a) If A, B are two finite-dimensional central simple algebras over k, so is $A \otimes_k B$.
b) Consider the set of finite-dimensional central simple algebras over k modulo Morita equivalence. This set is in bijection with central division rings over k of finite dimension. With the operation $[A] + [B] := [A \otimes_k B]$, this set forms an abelian group, called the **Brauer group** of k.

Lemma 13.9: If A is a finite-dimensional central simple algebra over k, then $A_e := A \otimes_k A^{\mathrm{op}} \simeq \operatorname{End}_k(A)$.

Proof. A is a simple algebra iff A is a simple A_e-module (A-bimodule). So $Z(A) = \operatorname{End}_{A_e}(A) \simeq k$ and A is finite-dimensional over k. Then by Theorem 3.3 (density theorem), $A_e \twoheadrightarrow \operatorname{End}_k(A)$. If $d = \dim_k(A)$, then $\dim_k(A_e) = \dim_k(\operatorname{End}(A)) = d^2$, so in fact this surjection is an isomorphism. □

Theorem 13.10 (Azumaya-Nakayama): Suppose A is a central simple algebra over k and B is any algebra over k. Then two-sided ideals in $A \otimes_k B$ are in bijection with two-sided ideals in B.

Proof. Our goal is to describe submodules of the $A_e \otimes_k B_e$-module $A \otimes_k B$. Consider $A \otimes_k B$ as an $A_e \otimes_k k$-module first. Then it's a simple module tensored with vector space. Hence A_e-submodules of it are of the form $A \otimes_k V$, $V \subset B$ a subspace (this follows from the classification of submodules in a semisimple module). But $A \otimes_k V$ is a $k \otimes_k B_e$-submodule iff V is a B_e-submodule of B, so in fact V must be a two-sided ideal of B. □

14.1 Definition and first properties of Brauer group

Lemma 14.1: The center of a simple ring is a field.

Proof. Saying that A is a simple ring, i.e. it has no nontrivial proper two-sided ideals, is equivalent to saying that A is simple as an $A \otimes_k A^{\mathrm{op}}$-module. Then $\mathrm{Hom}_{A\otimes_k A^{\mathrm{op}}}(A, A) = Z(A)$ and by Schur's Lemma 1.26, it must be a division ring. It remains to note that every commutative division ring is a field. □

Lemma 14.2: For A, B two algebras over k, $Z(A \otimes_k B) = Z(A) \otimes_k Z(B)$.

Proof. Suppose $x \in A \otimes B$ is central. We can write $x = \sum a_i \otimes b_i$ where the $a_i \in A$ are linearly independent and likewise for the $b_i \in B_i$. Then for all $a \in A$,

$$[x, a \otimes 1] = \sum [a, a_i] \otimes b_i = 0.$$

Since the b_i are linearly independent, this implies the a_i are all central. Likewise, $b_i \in Z(B)$. □

Proof (of Theorem 13.8).

a) By Theorem 13.10, the tensor product $A \otimes_k B$ is a simple ring, and by the above lemmas its center is the field $Z(A) \otimes_k Z(B) = k$.

b) The tensor operation is well-defined up to Morita equivalence since $A \sim \mathrm{Mat}_n(A)$ and

$$\mathrm{Mat}_n(A) \otimes_k B = \mathrm{Mat}_n(k) \otimes_k A \otimes_k B = \mathrm{Mat}_n(A \otimes_k B).$$

The operation is obviously commutative and associative, has identity k, and inverse $-[A] = [A^{\mathrm{op}}]$ since $[A \otimes_k A^{\mathrm{op}}] = [\mathrm{End}_k(A)] = [k]$.
To see that the set is in bijection with division rings over k of finite dimension, note that Theorem 2.17 implies that any central simple algebra A with center k has the form $\mathrm{Mat}_n(D)$ where D is a skew field with center k. D is unique because we can define D as $\mathrm{End}_A(L)^{\mathrm{op}}$ where $L \cong D^n$ is the unique simple A-module (see Artin-Wedderburn theorem 2.17).

□

Example 14.3: $\mathrm{Br}(\mathbb{R}) \cong \mathbb{Z}/2\mathbb{Z}$ because there are exactly two finite-dimensional skew fields over $\mathbb{R}$, namely $\mathbb{R}$ and $\mathbb{H}$.

Lemma 14.4: If E/F is a field extension, then A is a central simple algebra over F iff $A \otimes_F E$ is a central simple algebra over E. More generally, if B is an algebra over E, and A is an algebra over F, then $A \otimes_F B$ is a central simple algebra over E iff A is a central simple algebra over F and B is a central simple algebra over E.

Proof. Assume that A/F and B/E are central simple algebras. Then, by Theorem 13.10, $A \otimes_F B$ is a simple ring. Its center is (by lemma 14.2):

$$Z(A \otimes_F B) = Z(A) \otimes_F Z(B) = F \otimes_F E = E.$$

Assume now that $A \otimes_F B$ is a central simple algebra over E. Again from lemma 14.2 we know that $Z(A) \otimes_F Z(B) = Z(A \otimes_F B) = E$ so we must have $Z(A) = F$, $Z(B) = E$. It remains to note that if A is not simple, then there exists a nonzero proper two-sided ideal $I \subset A$ but then $I \otimes_F B$ will be a nonzero proper two-sided ideal in $A \otimes_F B$. Then A is simple, and ideals in $A \otimes_F B$ are in bijection with ideals in B by Azumaya-Nakayama theorem 13.10, so B is simple as well. □

Corollary 14.5: If E/F is a field extension, it induces a group homomorphism called the **base change map**

$$\mathrm{Br}(F) \to \mathrm{Br}(E), [A] \mapsto [A \otimes_F E].$$

Proof. It's a group homomorphism because

$$(A \otimes_F E) \otimes_E (E \otimes_F B) \cong E \otimes_F (A \otimes_F B).$$

□

Example 14.6: Algebraically closed fields have no finite skew field extensions, so if $k = \bar{k}$ then $\mathrm{Br}(k) = 0$. This implies that all central simple algebras over such k are of the form $\mathrm{Mat}_d(k)$.

Definition 14.7: Let A be a central simple algebra over an arbitrary field F. The **degree of** A is the d such that

$$A \otimes_F \bar{F} \cong \mathrm{Mat}_d(\bar{F}).$$

Alternately, it is the d such that $\dim_F(A) = d^2$.

Definition 14.8: The kernel of the base change map for an extension E/F is denoted $\mathrm{Br}(E/F)$.

Definition 14.9: Let A be a central simple algebra over F. We say an algebraic field extension E/F **splits** A, or that A **splits over** E, if $[A] \in \mathrm{Br}(E/F)$, i.e. $A \otimes_F E \cong \mathrm{Mat}_n(E)$.

Example 14.10: Every central simple algebra A over F will split over $\bar{F}$.

Corollary 14.11: Every central simple algebra A over F will split over a finite extension, namely the one generated by the matrix coefficients of the isomorphism $A \otimes_F \bar{F} \cong \mathrm{Mat}_n(\bar{F})$ (in some bases of A, $\mathrm{Mat}_n(F)$).

14.2 Torsors and Galois forms

Classifying the central simple algebras of a fixed degree over a fixed field F splitting over a fixed field extension of E is a special case of **Galois forms** or the **Galois descent problem**. Here is an overview of the general procedure and the classification:

Assume that E/F is Galois. Then consider the set I of all E-linear isomorphism $A \otimes_F E \cong \mathrm{Mat}_n(E)$. $\mathrm{PGL}_n(E)$ acts on $\mathrm{Mat}_n(E)$ by conjugation; in fact, it is isomorphic to the group of automorphisms of $\mathrm{Mat}_n(E)$ (either a special case of the Theorem 14.15, see below, or a direct computation).

Hence, $\mathrm{PGL}_n(E)$ acts on I by sending an isomorphism $A \otimes_F E \cong \mathrm{Mat}_n(E)$ to $A \otimes_F E \cong \mathrm{Mat}_n(E) \xrightarrow{\text{conj}} \mathrm{Mat}_n(E)$. It turns out that this action is *simply transitive*. On the other hand, we have an action of the Galois group $G = \mathrm{Gal}(E/F)$ on both $A \otimes_F E$ and on $\mathrm{Mat}_n(E)$, so it acts on I by conjugation. These actions of $\mathrm{PGL}_n(E)$ and G are compatible. This defines what we call a $\mathrm{PGL}_n(E)$**-torsor over** G.

Hence, to every central simple algebra A of degree d split over E, we can assign a corresponding $\mathrm{PGL}_d(E)$-torsor over G, and it is not hard to see that this is a bijection. For example, the trivial torsor, where $I = \mathrm{PGL}_n(E)$, corresponds to $A \cong \mathrm{Mat}_n(F)$.

We will see in the next lecture that isomorphism classes of such torsors are classified by the nonabelian cohomology group $H^1(G, \mathrm{PGL}_n(E))$.

Moreover, this method generalizes to other algebraic objects depending on the choice of the base field, as long as "base change under field extension" makes sense: fix a reference object S, then the objects whose base change to E are isomorphic to S are in bijection with $\mathrm{Aut}(S)$-torsors over G.

14.3 Centralizer of a commutative subfield

Lemma 14.12: If $k \subset F \subset A$ where A is a central simple algebra over a field k, F is a field, $\dim_k(A) = d^2$, $[F : k] = n$, and $B = Z_A(F)$, then $\dim_F(B) = \left(\frac{d}{n}\right)^2$ and B is a central simple algebra over F also. That is,

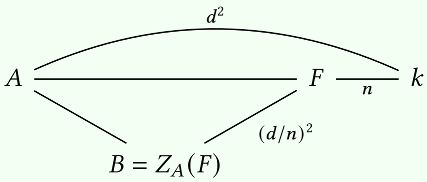

.

Moreover, $[B] = [A \otimes_k F] \in \mathrm{Br}(F)$.

In a sense, this is saying that in the chain of inclusions

$$k \subset F \subset Z_A(F) \subset A,$$

the index $[F : k]$ "mirrors" the "index" $[A : Z_A(F)]$, so that the larger F gets, the smaller $Z_A(F)$ gets.

Proof. $A \otimes_k F$ is a central simple algebra over F, and moreover it acts on A by $a \otimes f : x \mapsto axf$. So $\mathrm{End}_{A \otimes_k F}(A) = Z_A(F) = B$ is also a central simple algebra and is Morita equivalent to $A \otimes_k F$ (recall that we have the natural identification $A \otimes_k A^{\mathrm{op}} \xrightarrow{\sim} \mathrm{End}_k(A)$ and $Z_A(F) \otimes_k A^{\mathrm{op}}$ identifies with $\mathrm{End}_F(A) \subset \mathrm{End}_k(A)$ so, by Lemma 14.4, $B = Z_A(F)$ is indeed a c.s.a. over F).
To find $\dim_F(B)$, notice that for any central simple algebra C over F and a C-module M with $E = \mathrm{End}_C(M)$, we have

$$\dim_F(C) \dim_F(E) = \dim_F(M)^2.$$

Moreover, $C \otimes_F E \cong \mathrm{End}_F(M)$. This is because any simple algebra is a matrix algebra over a division ring, so in particular $C = \mathrm{Mat}_n(D)$, and $M = (D^n)^m$ is just a direct sum of the unique simple C-module (for some m) and $E = \mathrm{Mat}_m(D^{\mathrm{op}})$. Then

$$C \otimes_F E = \mathrm{Mat}_{nm}(D \otimes D^{\mathrm{op}}) = \mathrm{Mat}_{nmd}(F) = \mathrm{End}_F(M)$$

where $d = \dim_F(D)$, and taking dimensions we get the desired identity.
Setting $C = A \otimes_k F, M = A, B = E$, we get

$$n^2 \dim_F(B) = d^2 \Rightarrow \dim_F(B) = \left(\frac{d}{n}\right)^2.$$

□

Corollary 14.13: Let A be a central simple algebra of degree d over a field k. Then every subfield F of A has degree $\leqslant d$ over k. Moreover, field F is a maximal commutative subalgebra of A iff $[F : k] = d$.

Proof. The fact that $[F : k] \leqslant d$ directly follows from Lemma 14.12.
If $F \subset A$ is maximal commutative, then $Z_A(F)$ must be equal to F (indeed, otherwise there exists an element $x \in Z_A(F) \setminus F$ so $F[x]$ is a commutative subalgebra of A that is bigger than F). So $Z_A(F) = F$ and the claim about the dimension of F (over k) follows from Lemma 14.12. □

Warning 14.14: It may happen that $F \subset A$ is a maximal commutative *subfield* but not a maximal commutative *subalgebra* (take, for example, $A = \mathrm{Mat}_n(k)$ and $F = k$). If A is a skew field, then these two properties do coincide.

14.4 Skolem-Noether

Theorem 14.15 (Skolem-Noether): Let A be a simple Artinian ring with center k and B a simple finite-dimensional k-algebra. Then any two k-linear homomorphisms $B \longrightarrow A$ are conjugate by an invertible element of A.

This allows us to relate different embeddings of a given field in a central simple algebra.

Proof. Let $\varphi\colon B \longrightarrow A, \psi\colon B \longrightarrow A$ be two k-linear maps $B \longrightarrow A$. These give A two structures as an (A, B)-bimodule: A_φ where

$$a \otimes b\colon x \mapsto ax\varphi(b)$$

and A_ψ where

$$a \otimes b\colon x \mapsto ax\psi(b).$$

Since $A \otimes_k B^{\mathrm{op}}$ is simple (Theorem 13.10) and finitely generated as an A-module, it must be Artinian. So $A \otimes_k B^{\mathrm{op}}$ has only one simple module L, and any module M finitely generated over A will be isomorphic to L^n, $n < \infty$, and n is determined by the isomorphism class of $M|_A$. Then $A_\varphi \cong A_\psi$. The isomorphism is given by right multiplication by some left invertible, hence invertible, element of A that conjugates φ into ψ. □

14.5 Artin-Wedderburn

Theorem 14.16 (Artin-Wedderburn): There are no finite noncommutative skew fields. Hence, the Brauer group of a finite field is trivial.

Proof. Suppose that D is a noncommutative finite skew field with center $F = \mathbb{F}_q$. Let $E \subset D$ be a maximal commutative subfield. So by Corollary 14.13, $[E : F] = d$ where $d^2 = \dim_F(D)$. For $\alpha \in D$, $K = F(\alpha)$ will have degree d' over F with $d' \mid d$.
Then $E = \mathbb{F}_{q^d}$ and $K = \mathbb{F}_{q^{d'}}$. This implies that K is isomorphic to a subfield in E as an extension of F. This gives us two homomorphisms $E \longrightarrow D$ and $K \longrightarrow D$, so there exists an $x \in D^\times$ such that $xKx^{-1} \subset E$ by Theorem 14.15. $D^\times$ is a finite group and $E^\times \subset D^\times$ is a subgroup, and the following lemma implies that $E = D$.

Lemma 14.17: Let $H \subset G$ be a subgroup in a finite group G. If every element in G is conjugate to an element in H, then $H = G$.

Proof. Let $\mathcal{C}$ be the set of conjugacy classes in G. For each conjugacy class $C \in \mathcal{C}$, we know $|C| = |G : Z_G(g)|$, $g \in C$, and $Z_G(g)$ is the centralizer of g. By assumption $C \cap H$ is nonempty for every conjugacy class, and we can bound

$$|C \cap H| \geqslant [H : C_H(g)] \geqslant \frac{[G : Z_G(g)]}{[G : H]} = \frac{|C|}{[G : H]}.$$

with equality when $C \cap H$ is single H-conjugacy class (first equality) and $Z_G(g) \subset H$ (second equality). In particular, if $g = 1$, we will always get a strict inequality. Then

$$|H| = \sum |C \cap H| > \frac{\sum |C|}{[G : H]} = \frac{|G|}{[G : H]},$$

contradiction. □

□

15.1 Separable splitting fields

Theorem 15.1: For a finite Galois extension E/F, we have a natural isomorphism

$$\mathrm{Br}(E/F) = H^2(\mathrm{Gal}(E/F), E^\times).$$

To use this theorem, we want to say that every element splits over a finite Galois extension. In characteristic 0, every finite extension is contained in a finite Galois extension and we proved that every element splits over a finite extension. In general, a field extension is contained in a Galois extension iff it is separable.

Proposition 15.2: Every element in $\mathrm{Br}(F)$ splits over a finite separable extension (and hence over a finite Galois extension).

Proof. Let D be a skew field with center F (so a central simple algebra over F). It's enough to show that there exists a commutative subfield $E \subset D$ such that $E \supsetneq F$ and E/F is separable (since any csa A over F is a matrix algebra over some skew field D with center F, hence $[A] = [D] \in \mathrm{Br}(F)$); then we can consider instead the centralizer $D' = Z_D(E)$; since $[Z_D(E)] = [E \otimes_F D]$, we are done by induction on $\dim_F D$, use Lemma 14.12.
Suppose such an E does not exist. Then, by field extension theory, for every $x \in D$, there exists n such that $x^{p^n} \in F$.

Lemma 15.3: Let A be an $\mathbb{F}_p$-algebra. For $x \in A$, we have $\mathrm{ad}(x)^p = \mathrm{ad}(x^p)$, where $\mathrm{ad}(x) : y \mapsto xy - yx$.

Proof. If a, b are commuting elements in an $\mathbb{F}_p$-algebra, then $(a-b)^p = a^p - b^p$. Applying this to $a = L_x$ and $b = R_x$, where L_x is left multiplication by x and R_x is right multiplication by x, we see that $(a-b)^p = \mathrm{ad}(x)^p$ while $a^p - b^p = \mathrm{ad}(x^p)$. □

Now we have two ways to finish the argument.
The first uses Engel's Theorem (see 18.745): if $\mathfrak{g} \subset \mathfrak{gl}_n(F)$ is a subalgebra consisting of nilpotent matrices, then $\mathfrak{g}$ is nilpotent. Equivalently, it is contained in the algebra of strictly upper triangular matrices in some basis. The lemma implies that $\mathrm{ad}(x)$ is nilpotent for all x. Hence, the image of D in the Lie algebra $\mathrm{End}_F(D)$ (via the map $x \mapsto \mathrm{ad}(x)$) is nilpotent by Engel's Theorem. This contradicts that $D \otimes_F E \cong \mathrm{Mat}_n(E)$ for some E.
The second uses Jordan normal form. Pick $x \in D$ such that $x \notin F$ but $x^p \in F$. Let $E = F(x)$. Then $[E : F] = p$ and $\dim_F(Z_D(E)) = \frac{d^2}{p}$. By the lemma, $\mathrm{ad}(x)^p = 0$ where $\mathrm{ad}(x) : D \longrightarrow D$, and

$$\dim_F(\ker(\mathrm{ad}(x))) = \dim_F(Z_D(E)) = \frac{\dim_F(D)}{p}.$$

Therefore, the Jordan normal form of $\mathrm{ad}(x)$ must have d^2/p equal Jordan blocks of size $p > 1$. In particular, $\ker(\mathrm{ad}(x)) \subset \mathrm{im}(\mathrm{ad}(x))$. So if $x \in \ker(\mathrm{ad}(x))$, there exists y such that $[x, y] = x$. Then $\mathrm{ad}(-y)$ fixes x, so $\mathrm{ad}(-y)$ cannot be nilpotent, contradiction. □

15.2 Group cohomology

Let G be a group. Recall that a G-module is the same as a $\mathbb{Z}[G]$-module, and for such a G-module M, we define

$$H^i(M) := \mathrm{Ext}^i_{\mathbb{Z}[G]}(\mathbb{Z}, M)$$

where $\mathbb{Z}$ is the trivial $\mathbb{Z}[G]$-module. In other words, H^i is the ith derived functor of the functor of G-invariants. To compute this, you can also use the bar resolution, which is a resolution for any flat algebra over a commutative ring, in particular $\mathbb{Z}[G]$. This results in a complex where C^n consists of maps $f : G^n \longrightarrow M$ and the differential is

$$df(g_0, \dots, g_n) = g_0 f(g_1, \dots, g_n) + \sum_{i=0}^{n-1} (-1)^i f(\dots, g_i g_{i+1}, \dots) + (-1)^n f(g_0, \dots, g_{n-1}).$$

Example 15.4: In particular, a 1-cocycle is a map $c\colon G \to M$ such that $gc(h) - c(gh) + c(g) = 0$; these are called "cross homomorphisms" and you can produce them from an M-torsor T over G and a choice of point $x_0 \in T$. The correspondence takes a cocycle c to the G-module structure on M where $g.m = m + c(g)$ ($T = M$ and $x_0 = 0$). Given a torsor T and a point $x_0 \in T$, for each $g \in G$ we set $c(g)$ to be the element in M such that $g(x_0) = x_0 + c(g)$. Varying the choice of a point results in adding a coboundary to the cocycle. We end up with a bijection between $H^1(G, M)$ and isomorphism classes of M-torsors over G. There is also a bijection between $H^1(G, M)$ and extensions of $\mathbb{Z}$ by M because of its definition as Ext^1.

Remark 15.5: Moreover, the definition of $H^1(G, M)$ generalizes to the case when M is a nonabelian group equipped with a G-action, and in this case we view M as acting on itself on the right, while G acts on the left. This does not hold for higher cohomology.

Example 15.6: A 2-cocycle is a map $c\colon G^2 \to M$ such that $gc(h,k) - c(gh,k) + c(g,hk) - c(g,h) = 0$.

Definition 15.7: A **cross-product extension of** G **by** M is a group $\tilde{G}$ with a normal subgroup identified with M and an isomorphism $\tilde{G}/M \cong G$ (i.e. an extension of G by M) such that the conjugation action of $\tilde{G}$ on M, which automatically factors through G, coincides with the given action of G on M (the cross-product).

2-cocycles are in bijection with cross-product extensions of G by M together with a splitting of the surjection of sets $\tilde{G} \to G$. Choosing a different splitting modifies the cocycle by a coboundary. Hence, there is a bijection between $H^2(G, M)$ and cross-product extensions of G by M up to isomorphism.

15.3 Cross-product algebras

Given a group G acting on a ring R, we can form the smash product

$$G\#R = \bigoplus_{g\in G} R_g, \quad x_g y_h = (xg(y))_{gh}.$$

Given a cocycle $c \in H^2(G, R^\times)$, one can define a twisted version of this called the **cross-product algebra**,

$$G\#_c R = \bigoplus_{g\in G} R_g, \quad x_g y_h = (xg(y)\underbrace{c(g,h)}_{\in R^\times})_{gh}.$$

Up to isomorphism, the cross-product algebra depends only on the class of c in $H^2(G, R^\times)$.

This can also be described in terms of the cross-product group $\tilde{G}$ as

$$\tilde{G}\#R/(\lambda - [\lambda]), \lambda \in R^\times, [\lambda] \in \tilde{G} \text{ is the corresponding element.}$$

16 April 11 - Cohomological description of the Brauer group

More on the Brauer group.

16.1 Cross-product algebras and Galois extensions

Proposition 16.1: Suppose E/F is a Galois extension. Then we have a bijection between central simple algebras over F with maximal commutative subfield (isomorphic to) E and cross-product extensions of $G = \mathrm{Gal}(E/F)$ by $E^\times$.

Proof. The bijection will send a central simple algebra A with maximal commutative subfield E to $\tilde{G} = \mathrm{Nm}_{A^\times}(E)$, where Nm is for normalizer; this is a cross-product extension of G by $E^\times$. Since conjugating by an element of $\tilde{G}$ induces a Galois automorphism of E by definition, there is a homomorphism $\tilde{G} \to G$. Skolem-Noether 14.15 implies that this is onto. The kernel of this homomorphism is the invertible elements of A that commute with E.

Since $Z_A(E) = E$, the kernel must be $E^\times$ and we have an exact sequence $0 \to E^\times \to \tilde{G} \to G \to 0$, giving us a cross-product extension.
In the other direction, the bijection will take a cross-product extension, which corresponds to $c \in H^2(G, E^\times)$, to $A := G\#_c E$. First, we claim that A is a central simple algebra. First, it is simple. Notice that $E \otimes_F E \cong \prod_G E$ (by Galois theory) and A is a free rank 1 module over $E \otimes_F E$. Conjugation by an element $x_g \in A_c$, $x \neq 0$, will permute the copies of E and send E_h to $E_{h'}$. Therefore, for a nonzero ideal $I \subset A$, I must have a nonzero intersection with some E_g, hence it contains E_g, but then I contains all the E_g and $I = A$.
And $Z_A(E) = E$: if $x = (x_g) \in A$ with $x_g \neq 0$ and $g \neq 1$, we can pick $y \in E$ such that $g(y) \neq y$, in which case

$$(xy)_g = g(y)x \neq yx = (yx)_g.$$

Hence, $Z(A) \subset E$ and $Z(A) = E^G = F$.
Now we check these are inverse bijections. Start with $\tilde{G} = \tilde{G}_c$ and let $A = G\#_c E$. Then $\mathrm{Nm}_{A^\times}(E) = \tilde{G}$, since if $a \in A^\times$ normalizes E, then $ag^{-1} \in Z(E)$ for some $g \in G$, so $ag^{-1} \in E^\times$. Conversely, starting with A, mapping to a cocycle c, the map $(x_g) \mapsto \sum xg$ is a homomorphism. Then the map $G\#_c E \to A$ is injective because $G\#_c E$ is simple, and moreover, these have the same dimension over F, so the map is an isomorphism. □

Remark 16.2: While the above gives a transparent relation between central simple algebras and cross-products, some questions about this construction turn out to be quite hard. In particular, it's hard to determine whether a given cross-product algebra is a skew field or whether a given skew field is isomorphic to a cross-product algebra, see e.g. [2].

16.2 Maximal commutative subfields and splitting fields

Lemma 16.3: Let E/F be a finite extension and A a central simple algebra over F. Then $[A] \in \mathrm{Br}(E/F)$ iff A is equivalent to an algebra A' containing E as a maximal subfield.

Proof. Suppose that $E \subset A$ is a maximal subfield. We know that $A \simeq \mathrm{Mat}_n(D)$ for some skew field D, so $[A] = [\mathrm{Mat}_n(D)] = [D] \in \mathrm{Br}(F)$. This means that (base changing to E) $[A \otimes_F E] = [D \otimes_F E] = Z_D(E)] = [E] = 0 \in \mathrm{Br}(E)$ (use the result from last lecture), hence $[A] = [D] \in \mathrm{Br}(E/F)$.
In the other direction, suppose that A splits over E. Write $A = \mathrm{Mat}_m(D)$, then consider the minimal n such that $A' = \mathrm{Mat}_n(D)$ contains E as a maximal subfield. Since $[A] = [A'] = [D] \in \mathrm{Br}(F)$, we have (also from last time) $[Z_{A'}(E)] = [A' \otimes_F E] = [A \otimes_F E] = 0 \in \mathrm{Br}(E)$. Now $Z_{A'}(E)$ is a skew field which is Morita equivalent to 0 in $\mathrm{Br}(E)$, hence is E itself. This implies that A' (which is Morita equivalent to A) contains E as a maximal subfield.
□

16.3 Proof of the theorem

Corollary 16.4: Let E/F be a finite Galois extension. Then $\mathrm{Br}(E/F) \cong H^2(\mathrm{Gal}(E/F), E^\times)$.

Proof. Now we know that there is a bijection between central simple algebras over F with maximal commutative subfield E and $H^2(G, E^\times)$. The Lemma 16.3 implies that every class $[A]$ in $\mathrm{Br}(E/F)$ has a representative A' with maximal commutative subfield E, hence there is a surjective map $\mathrm{Br}(E/F) \twoheadrightarrow H^2(G, E^\times)$ (currently just a map of sets, not a homomorphism). On the other hand, it is also an injection (on sets): Lemma 16.3 tells us there is an injection from $\mathrm{Br}(E/F)$ to c.s.a.s with maximal commutative subfield E, which are then in bijection with classes in $H^2(G, E^\times)$. So we have a bijection between $\mathrm{Br}(E/F)$ and $H^2(G, E^\times)$.
We need to check that this is a group homomorphism. Let's rewrite the group structure on H^2 in terms of cross-products. Given $\tilde{G}_{c_1}$ and $\tilde{G}_{c_2}$, one can check that

$$\tilde{G}_{c_1 c_2} \cong \tilde{G}_{c_1} \times_G \tilde{G}_{c_2} / (m, -m) \subset M \times M.$$

Now we want to check that

$$B := A_{c_1} \otimes_F A_{c_2} \sim A_{c_1 c_2}.$$

But $B \supset E \otimes_F E = \prod_G E$. Let $e = 1_1 \in E \otimes_F E$. Then $eBe \cong A_{c_1 c_2}$ and this represents the class $[A_{c_1}] + [A_{c_2}]$, so the group structures on both are compatible. ◻

16.4 Applications

Proof (of Theorem 14.16). Recall that we want to prove that there are no finite noncommutative skew fields. This is equivalent to proving that $\mathrm{Br}(\mathbb{F}_{q^n}/\mathbb{F}_q)$ is trivial for all n, i.e. by the above, that $H^2(G, \mathbb{F}_{q^n}^\times) = 0$. The Galois group of this extension is $\mathbb{Z}/n\mathbb{Z}$. Pick a generator $\gamma \in \mathbb{Z}/n\mathbb{Z}$. For cyclic groups, we can use the following resolution of $\mathbb{Z}$ to compute $H^*(\mathbb{Z}/n\mathbb{Z}, M)$:

$$\cdots \longrightarrow \mathbb{Z}[G] \xrightarrow{1+\cdots+\gamma^{n-1}} \mathbb{Z}[G] \xrightarrow{1-\gamma} \mathbb{Z}[G] \xrightarrow{1+\gamma+\cdots+\gamma^{n-1}} \mathbb{Z}[G] \xrightarrow{1-\gamma} \mathbb{Z}[G] \longrightarrow \mathbb{Z}$$

where the leftmost arrow fits in the exact sequence because

$$(1-\gamma)\sum_{i=0}^{n-1} n_i\gamma^i = \sum_{i=0}^{n-1}(n_i - n_{i-1})\gamma^i = 0 \Leftrightarrow n_i = n_j \ \forall i, j.$$

The complex is 2-periodic, since

$$\left(\sum_{i=0}^{n-1}\gamma^i\right)\left(\sum_{i=0}^{n-1} n_i\gamma^i\right) = \left(\sum_{i=0}^{n-1} n_i\right)\left(\sum_{i=0}^{n-1}\gamma^i\right).$$

So $H^{2k}(\mathbb{Z}/n\mathbb{Z}, M) = M^G/\mathrm{Im}(\mathrm{Av})$, where $\mathrm{Av}\colon M \longrightarrow M^G$ takes $m \mapsto \sum_{g\in G} g(m)$. Thus if E/F is a Galois extension with $G \cong \mathbb{Z}/n\mathbb{Z}$, which is our case, $\mathrm{Br}(E/F) = H^2(G, E^\times) = F^\times/\mathrm{Nm}(E^\times)$ where here Nm is the image of the norm map. But for $F = \mathbb{F}_q, E = \mathbb{F}_{q^n}$, $\mathrm{Nm}(x) = x^{(q^n-1)/q-1}$, so cyclicity of $E^\times$ implies that $\mathrm{Nm}\colon E^\times \twoheadrightarrow F^\times$ and $\mathrm{Br}(E/F) = 0$. ◻

Remark 16.5: We have shown that if E/F is a Galois extension with $G \cong \mathbb{Z}/n\mathbb{Z}$, then $\mathrm{Br}(E/F) = H^2(G, E^\times) = F^\times/\mathrm{Nm}(E^\times)$. The identification can be explicitly described as follows: recall that $\gamma \in G$ is a generator. Consider the "twisted polynomial algebra" $E\langle x;\gamma\rangle := \{\sum_i c_i x^i \mid c_i \in E\}$ with $xc = \gamma(c)x$ for $c \in E$. Pick $a \in F^\times$, the corresponding central simple algebra is $E\langle x;\gamma\rangle/(x^n - a)$ (such algebras are called *cyclic* algebras).

Example 16.6: $\mathrm{Br}(\mathbb{C}/\mathbb{R}) = \mathbb{R}^\times/\mathrm{Nm}(\mathbb{C}^\times) = \mathbb{Z}/2\mathbb{Z}$. It is easy to see that the element $[1] \in \mathbb{Z}/2\mathbb{Z}$ corresponds to the central simple algebra $\mathbb{H}$ of quaternions.

16.5 Index and period

Definition 16.7: The **index** of an element in a Brauer group is the degree of its minimal representative. That is, the index of $[\mathrm{Mat}_n(D)] = [D]$ equals d if D is a skew field of dimension d^2.

Definition 16.8: The **period** of a central simple algebra A over F is the order of $[A] \in \mathrm{Br}(F)$.

Lemma 16.9: The period of an element in the Brauer group divides its index. In particular, the period is always finite, and Br is torsion.

Proof. Let D be the skew field representative of this element, say it has degree d, with center F. We proved that D has a maximal subfield E such that E/F is separable in Proposition 15.2. Let K be a Galois extension of F containing E and $G = \mathrm{Gal}(E/F)$. Then $E = K^H$ for an index d subgroup $H \subset G$, $H = \mathrm{Gal}(K/E)$.
Now the lemma follows from the following fact about group cohomology: given a finite group G, $H \subset G$ of index d, and a G-module M, the kernel of $\mathrm{res}\colon H^i(G, M) \longrightarrow H^i(H, M)$ is killed by d. This is because we can define a map $a\colon H^i(H, M) \longrightarrow H^i(G, M)$ so that $a \circ \mathrm{res}$ is multiplication by d. For $i = 0$, this map sends $m \mapsto \sum_{g\in G/H} g(m)$, and in higher degrees, take an injective resolution of M over G, which will restrict to an injective resolution over H, then apply the above map to each term of the resolution.
Hence, the dth power of every element in the Brauer group vanishes. ◻

Not all integers arise as indexes of Brauer classes:

Lemma 16.10: If F is a perfect characteristic p field, the Brauer group has no p-torsion.

Proof. A separable finite extension E of F is also perfect. Hence $E^\times \to E^\times, x \mapsto x^p$ is an isomorphism, so it induces an automorphism $H^2(G, E^\times) \to H^2(G, E^\times)$. □

Finally, we give a cohomological description of $\mathrm{Br}(F)$ in terms of the absolute Galois group. We can describe by taking a limit of the $\mathrm{Br}(E/F)$, but we need to take into account that the absolute Galois group $G_F = \mathrm{Gal}(\bar{F}_{\mathrm{sep}}/F)$ (where $\bar{F}_{\mathrm{sep}}$ is the separable algebraic closure) is a profinite group. Hence, we need to consider continuous cohomology instead of normal cohomology, where all cocycles in the standard complex are required to be continuous. Then we can show that

$$H^2_{\mathrm{cont}}(G_F, \bar{F}_{\mathrm{sep}}^\times) = \varinjlim_E H^2(\mathrm{Gal}(E/F), E^\times) = \varinjlim \mathrm{Br}(E/F) = \mathrm{Br}(F).$$

17.1 Reduced norm and trace

We can generalize the determinant and trace to central simple algebras. Suppose A is a central simple algebra of degree d over k.

Proposition 17.1: There exist unique polynomial maps $\tau, \delta \colon A \to k$ so that for any field extension K/k such that A splits over K,

$$\tau_K \colon A \otimes_k K \cong \mathrm{Mat}_n(K) \to K$$

is the trace and

$$\delta_K \colon A \otimes_k K \cong \mathrm{Mat}_n(K) \to K$$

is the determinant. τ is called the **reduced trace** and δ is called the **reduced norm**.

Example 17.2: Let's take $A = \mathbb{H}$ and $k = \mathbb{R}$. Then $\tau \colon a+bi+cj+dk \mapsto 2a$ and $\delta \colon a+bi+cj+dk \mapsto a^2+b^2+c^2+d^2$.

Proof. By the Artin-Wedderburn theorem, WLOG we can assume $|k| = \infty$ so that we can say that polynomials are determined by their values on k^n. Now the proof follows from Galois descent and the fact that Tr, det are invariant under all automorphisms of the matrix ring. For a fixed extension K/k, τ, δ satisfying the compatibility with Tr, det are unique; moreover, they will satisfy the same compatibility for any extension $K' \supset K$, and also for $K'' \subset K$ if K splits A. So we only have to construct τ, δ satisfying the compatibility for a fixed extension splitting A.
Choose a finite Galois extension K/k which splits A and choose an isomorphism $A \otimes K \cong \mathrm{Mat}_n(K)$. Let $G = \mathrm{Gal}(K/k)$, it acts on $A \otimes K$ by acting on K. It suffices for us to show that det, Tr commute with the G-action, which will imply that they come from polynomial maps defined over k.
To see this, consider the action of G on $\mathrm{Mat}_n(K)$, which is different from the action above; say it sends $a \mapsto {}^\gamma a$. Then the map $a \mapsto \gamma^{-1}({}^\gamma a)$ is a K-linear automorphism on $\mathrm{Mat}_n(K)$, hence given by conjugation by some element $g_\gamma \in \mathrm{GL}_n(K)$. Since det is conjugation-invariant, we have

$$\det(a) = \det(\gamma^{-1}({}^\gamma a)) \Rightarrow \det(\gamma(a)) = \det({}^\gamma a) = \gamma(\det a).$$

The same argument works for trace. So we are done. □

From these, we see that $\tau(ab) = \tau(ba)$, $\delta(ab) = \delta(a)\delta(b)$, and $\delta(1) = 1$.

17.2 C_1 fields

Definition 17.3: We say a field is **quasi-closed** or C_1 if any homogeneous polynomial of degree d in $n > d$ variables has a nontrivial zero. More generally, we say a field is C_k if any homogeneous polynomial of degree d in $n > d^k$ variables has a nontrivial zero.

Proposition 17.4: If F is C_1, $\mathrm{Br}(F) = 0$.

Proof. Suppose not. Then let D be a skew field finite over F with $Z(D) = F$. Then δ (the reduced norm) is a degree d polynomial but $\dim_F(D) = d^2$, so δ has a nontrivial zero. But $\delta(D^\times) \subset F^\times$ is invertible, a contradiction since we just said that δ has a nontrivial zero, i.e. a zero in $D - \{0\} = D^\times$. □

Lemma 17.5: Finite extensions of C_1 fields are also C_1.

Proof. Suppose F is C_1 and E/F is a degree m extension. Let P be a polynomial of degree d in n variables over E. By choosing a basis for E over F, we can identify $E^n = F^{nm}$. Then consider the polynomial

$$\tilde{P}(x) := \mathrm{Nm}_{E/F}(P(x));$$

this is a degree md polynomial in mn variables over F, and it has a nontrivial zero iff P does. □

Theorem 17.6 (Chevalley-Warning): Finite fields are C_1 fields.

Proof. The previous lemma shows that it's enough to consider $\mathbb{F}_p$. Then the result follows from the following fact: if P is a homogeneous polynomial in n variables of degree $n > d$ over $\mathbb{F}_p$, the number of zeroes is 0 mod p. Since there is at least one zero (the trivial one), there are at least p zeroes. So it remains to prove this fact.
We know that for $a \in \mathbb{F}_p$, a^{p-1} is either 0 or 1 (if $a \neq 0$). So

$$\sum_{a_1,\ldots,a_n \in \mathbb{F}_p} (1 - P(a_1, \ldots, a_n)^{p-1}) \equiv \#\,\text{zeroes of } P \pmod{p}.$$

Every monomial in this sum (considered as a polynomial in a_i) will have at least one variable that has exponent less than $p - 1$ because the polynomial has degree $d(p - 1)$ and has n variables (we use that $d(p - 1) < n(p - 1)$ because $d < n$). Summing over that variable and using that $\sum_a a^m = 0$ when $0 \leqslant m < p - 1$, we see that the whole sum is 0. □

Remark 17.7: This gives another proof of Theorem 14.16.

Theorem 17.8 (Tsen's Theorem): Suppose k is algebraically closed. Then the field $F = k(t)$ is C_1.

Proof (Sketch). Clear denominators so that WLOG $P \in k[t][x_1, \ldots, x_n]$. Then use that a system of m homogeneous polynomial equations over k in n variables has a nontrivial solution if $n > m$ (this is true because k is algebraically closed). If K is the maximum degree (in t) of a coefficient of P, look at a solution of degree r. Then you get $dr+K+1$ equations in $(r + 1)n$ variables and $d < n$ implies $dr + K + 1 < (r + 1)n$ when $r \gg 0$. □

17.3 Second approach to the cohomological description of Brauer group

Let A be a central simple algebra over F and E/F a finite Galois extension. As described in the proof of Proposition 17.1, when you fix an isomorphism $A \otimes_F E \cong \mathrm{Mat}_n(E)$, you get two G-actions, $\gamma(a)$ and ${}^\gamma a$, that differ by conjugation by $g_\gamma \in \mathrm{GL}_n(E)$. This g_γ is determined up to multiplication by a scalar matrix, so $g_{\gamma_1} g_{\gamma_2}$ and $g_{\gamma_1\gamma_2}$ have the same image in $\mathrm{PGL}_n(E) = \mathrm{Aut}(\mathrm{Mat}_n(E))$ (but lifting to GL_n requires a choice). So we can define

$$c(\gamma_1, \gamma_2) = g_{\gamma_1} g_{\gamma_2} g_{\gamma_1\gamma_2}^{-1} \in E^\times.$$

In fact, c is a 2-cocycle, and its class in H^2 is independent of choice. Therefore, we get a map $\mathrm{Br}(E/F) \to H^2(G, E^\times)$, and it's an isomorphism.

Remark 17.9: We can interpret the definition of c as follows. The set of isomorphisms $A \otimes_F E \cong \mathrm{Mat}_n(E)$ form a $\mathrm{PGL}_n(E)$-torsor over G. As discussed earlier, the isomorphism class of this torsor corresponds to an element $\tilde{c} \in H^1(G, \mathrm{PGL}_n(E))$, the nonabelian cohomology group. A short exact sequence of abelian groups with a G-action will produce a long exact sequence in cohomology. For

$$1 \to E^\times \to \mathrm{GL}_n(E) \to \mathrm{PGL}_n(E) \to 1$$

the first few terms of the sequence are still well-defined, even though the sequence involves two nonabelian groups. The class c is the image of $\tilde{c}$ under the connecting homomorphism.
The injectivity of the map can be deduced from Hilbert's Theorem 90, which says that $H^1(G, \mathrm{GL}_n(E)) = 1$. (Hilbert originally considered the case $n = 1$ only.) An equivalent form of this statement is as follows: given an n-dimensional E-vector space V_E with a compatible G-action, there is an F-vector space V_F and a G-equivariant isomorphism $V_E = V_F \otimes_F E$.

17.4 Brauer groups of local fields

Theorem 17.10: Let F be a non-Archimedean local field, i.e. it's a finite extension of $\mathbb{Q}_p$ or $\mathbb{F}_p((t))$ (in which case $F \cong \mathbb{F}_q((t))$). Then $\mathrm{Br}(F) \cong \mathbb{Q}/\mathbb{Z}$.

First, let us recall without proof some facts about non-Archimedean local fields. If F is such a field, we have a valuation $F^\times \to \mathbb{Z}$ satisfying $v(ab) = v(a) + v(b)$ and $v(a+b) \geqslant \min(v(a), v(b))$; we can extend this to F by setting $v(0) = \infty$. WLOG we can assume that v is onto. Then there exists an element π with $v(\pi) = 1$, called a uniformizer. The elements x with $v(x) \geqslant 0$ form the ring of integers $\mathcal{O} \subset F$, the elements x with $v(x) \geqslant 1$ form the unique maximal ideal $\mathfrak{m} = \pi\mathcal{O} \subset \mathcal{O}$, and the residue field $k = \mathcal{O}/\pi\mathcal{O}$ is finite. For all $x \in F^\times$, $x\pi^{-v(x)} \in \mathcal{O}^\times$.

Definition 17.11: If E/F is a finite extension, then k_E/k_F is an extension of finite fields. Its degree $i_{E/F} = [k_E : k_F]$ is the **inertia degree** of the extension. The **ramification index** of the extension, $r = r_{E/F}$, is the integer such that $\pi_E^r \pi_F^{-1} \in \mathcal{O}^\times$ where π_E, π_F are uniformizers of their respective valuations. Then

$$[E : F] = i_{E/F} r_{E/F}$$

since you can see these are both $\dim_{k_F}(\mathcal{O}_E/\mathfrak{m}_E)$.

Remark 17.12: This also works if E is a skew field.

Definition 17.13: If $r = 1$, we say that E/F is **unramified**. In this case, E/F is Galois and $\mathrm{Gal}(E/F) \cong \mathrm{Gal}(k_E/k_F)$ (in particular, it is cyclic).

Proposition 17.14: Every central simple algebra over a local field F splits over an unramified extension.

Proof (Sketch). Let D be a central simple algebra over F. Then we can extend the valuation to $D^\times$, choose a uniformizer π_D where $v_D(\pi_D) = 1$, $\mathcal{O}_D = \{x \in D \mid v_D(x) \geqslant 0\}$. We get a finite extension $k_D := \mathcal{O}_D/\pi_D\mathcal{O}_D$ over k_F (note that by Artin-Wedderburn theorem, k_D is a field), and

$$\dim_F D = d^2 = [k_D : k_F] r_{D/F}$$

where d is the degree of D. We also claim that $i_{D/F}, r_{D/F} \leqslant d$ (recall that $i_{D/F} := [k_D : k_F]$). To see this, it's enough to show the existence of commutative subfields E_1, E_2 in D with $i_{D/F} \leqslant [E_1 : F]$ and $r_{D/F} \leqslant [E_2 : F]$ (use Corollary 14.13). Let $E_1 = F(\alpha)$ where $\alpha \in \mathcal{O}_D$ is such that $\alpha \bmod \pi_D\mathcal{O}_D$ generates k_D over k_F and $E_2 = F(\pi_D)$.
Therefore, $i_{D/F} = r_{D/F} = d = [E_1 : F]$. This shows that E_1/F is unramified and that it is a maximal commutative subfield in D. Thus it splits D (see Lemma 16.3) and is our desired extension. □

Proposition 17.15: If E/F is an unramified degree n extension of a non-Archimedean local field, then $\mathrm{Br}(E/F) = \mathbb{Z}/n\mathbb{Z}$.

Proof. We saw last time that for a cyclic extension, $\mathrm{Br}(E/F) \cong F^\times/\mathrm{Nm}(E^\times)$. Since E/F is unramified, $\mathrm{Gal}(E/F) \cong \mathrm{Gal}(k_E/k_F)$ and every extension of finite fields is cyclic (the Galois group is generated by the Frobenius). For an unramified extension, $O_E^\times \twoheadrightarrow O_F^\times$; this follows from surjectivity of the associated graded maps $k_E^\times \twoheadrightarrow k_F^\times$ and $(1+\pi^n O_E)/(1+\pi^{n+1}O_E) \twoheadrightarrow (1+\pi^n O_F)/(1+\pi^{n+1}O_F)$, where $\pi = \pi_F$. The first map is identified with the norm and the second with the trace $k_E \longrightarrow k_F$. Since $\mathrm{Nm}(\pi) = \pi^n$, we get that $\mathrm{Br}(E/F) = \mathbb{Z}/n\mathbb{Z}$. □

Proof (of Theorem 17.10). Let F^{unr} be a maximal unramified extension of F. Then it contains a unique degree n subextension F_n/F for every $n > 1$ and

$$\mathrm{Br}(F) = \mathrm{Br}(F^{\mathrm{unr}}/F) = \varinjlim \mathrm{Br}(F_n/F) = \varinjlim \mathbb{Z}/n\mathbb{Z} = \mathbb{Q}/\mathbb{Z}.$$

□

Remark 17.16: The theorem allows us to formulate a version of the reciprocity law of Class Field Theory. Let k be a global field, i.e. a finite extension of $\mathbb{Q}$ or $\mathbb{F}_p(t)$. For every valuation v, we get a corresponding local field k_v by completing k at v. Then we get a map

$$\mathrm{Br}(k) \to \prod_v \mathrm{Br}(k_v)$$

and we claim that in fact

$$\mathrm{Br}(k) \hookrightarrow \bigoplus_v \mathrm{Br}(k_v)$$

and this induces an isomorphism of $\mathrm{Br}(k)$ with the kernel of the sum map, i.e.

$$\mathrm{Br}(k) \cong \left\{ (b_v) \in \bigoplus_v \mathrm{Br}(k_v) \mid \sum b_v = 0 \right\} = \ker\left(\bigoplus_v \mathrm{Br}(k_v) \to \mathbb{Q}/\mathbb{Z} \right).$$

This is one of several equivalent forms of the reciprocity law of class field theory. For example, the corresponding identity for degree 2 central simple algebras over $\mathbb{Q}$, $\mathbb{H}_{a,b} = \mathbb{Q}\langle i, j\rangle/(i^2 = a, j^2 = b, ij = -ji)$ is essentially equivalent to quadratic reciprocity.

18.1 Azumaya algebras

Let R be a commutative ring and A a ring over R that is **finitely generated and projective** (equivalently, locally free) as an R-module. Then the rank is a locally constant function on $\mathrm{Spec}(R)$. Assume this function is nowhere vanishing. It will also be occasionally convenient for us to assume that the rank is constant. Let us use the notation $A_S := S \otimes_R A$ for a homomorphism of rings $R \longrightarrow S$.

Lemma 18.1: For R, A as above the following are equivalent:

a) The map $A \otimes_R A^{\mathrm{op}} \longrightarrow \mathrm{End}_R(A)$ is an isomorphism.

b) For every algebraically closed field k and a homomorphism $R \longrightarrow k$, $A_k \cong \mathrm{Mat}_n(k)$.

c) For every maximal ideal $\mathfrak{m} \subset R$, let $k = R/\mathfrak{m}$; the ring A_k is a central simple algebra over k.

Proof. We check that *a)* $\Rightarrow$ *b)* $\Rightarrow$ *c)* $\Rightarrow$ *a)*.
a) $\Rightarrow$ *b)*: since A is locally free, for every $R \longrightarrow S$, $\mathrm{End}_S(A_S) = (\mathrm{End}_R(A))_S$. Thus property *a)* is inherited by base change and A_k is a finite-dimensional central simple k-algebra. Hence it's isomorphic to $\mathrm{Mat}_n(k)$ when k is algebraically closed.
b) $\Rightarrow$ *c)*: since $A_k \otimes_k \bar{k} \cong \mathrm{Mat}_n(\bar{k})$, A_k is a central simple algebra.

$c) \Rightarrow a)$: If $\varphi\colon M \to N$ is a map of finitely generated modules over a commutative ring where N is projective that induces an isomorphism $M_k \to N_k$ for every $k = R/\mathfrak{m}$, then it is an isomorphism. This is because Nakayama's Lemma implies φ is surjective, so N projective implies $M \cong N \oplus \ker\varphi$; then another application of Nakayama's Lemma shows that $\ker\varphi = 0$. Applying this fact to $M = A \otimes_R A^{\mathrm{op}}$, $N = \mathrm{End}_R(A)$, we get the claim. □

Definition 18.2: A ring A satisfying the equivalent conditions of the lemma is called an **Azumaya algebra** over R.

We can think of an Azumaya algebra as a "globalized" version of central simple algebra, made more precise by part (c) of the lemma.

Example 18.3: Let R be a Noetherian domain and A an R-algebra finitely generated as an R-module. Let F be the field of fractions of R, and suppose A_F is a central simple algebra over F. Then, for a finite localization $S = R_{(r)}$ ($r \in R$), the ring A_S is an Azumaya algebra over S.

Example 18.4 (Differential operators in char p): Let k be a characteristic p field and $A = k\langle x, y\rangle/(yx - xy - 1)$ be the Weyl algebra. Then x^p, y^p are central in A since $\mathrm{ad}(x^p) = \mathrm{ad}(x)^p$ while $\mathrm{ad}(x)^2(y) = [x, 1] = 0$, likewise for y. We claim that A is an Azumaya algebra over $R = k[x^p, y^p]$.
One can check that $\{x^m y^n \mid m, n \in \mathbb{Z}_{\geqslant 0}\}$ form a k-basis in A, so A is a free module over $k[x^p, y^p]$ with basis $x^m y^n$, $m, n \in \{0, \ldots, p-1\}$. To check that it's an Azumaya algebra, it suffices to check this holds after an extension of scalars to $\bar{k}$, so WLOG we can assume that $k = \bar{k}$. Then by the Hilbert Nullstellensatz, maximal ideals in R are generated by $x^p - a, y^p - b$ for $a, b \in k$. Then

$$A_{a,b} := A/(x^p - a, y^p - b) \cong \mathrm{Mat}_p(k) = \mathrm{End}_k(k[x]/x^p)$$

where the isomorphism sends x to "multiplication by $x + \alpha$" and y to $\frac{\partial}{\partial x} + \beta$, where $\alpha^p = a$, $\beta^p = b$.

Example 18.5 (Quantum torus): Let $A = \mathbb{C}\langle z, z^{-1}, t, t^{-1}\rangle/(zt = qtz)$ for $q \in \mathbb{C}^\times$ fixed constant. If q is a primitive order ℓ root of unity, then A is an Azumaya algebra over $\mathbb{C}[z^\ell, z^{-\ell}, t^\ell, t^{-\ell}]$; the proof is similar to the previous example.

Azumaya algebras allow us to define the notion of a Brauer group over a ring. In particular, if A, B are Azumaya algebras over R, then so is $A \otimes_R B$.

Definition 18.6: The **Brauer group** of a ring R is the set of Morita equivalence classes of Azumaya algebras over R; $[A] + [B] = [A \otimes_R B]$, $-[A] = [A^{\mathrm{op}}]$.

Note that $[A^{\mathrm{op}}] = -[A]$ because $A \otimes_R A^{\mathrm{op}} \simeq \mathrm{End}_R(A)$, and R is Morita equivalent to $\mathrm{End}_R(A)$ iff A is a finitely generated projective generator. But by assumption A is a finitely generated projective, and the generator part follows from the fact that A restricted to every closed point of $\mathrm{Spec}\, R$ is a central simple algebra, hence nonzero.

For a homomorphism $R \to S$, we have a base change homomorphism $\mathrm{Br}(R) \to \mathrm{Br}(S)$, given by $A \mapsto A_S$; this is a homomorphism since A_S is Azumaya over S and $(A \otimes_R B)_S \cong A_S \otimes_S B_S$.

Remark 18.7: $[A] = 0$ iff $A \cong \mathrm{End}_R(M)$ where M is a finitely generated projective constant rank module over R. This means that M restricted to every closed point $\mathfrak{m}$ should be k^n for some fixed n, so $A|_{\mathfrak{m}} \simeq \mathrm{End}_{R_{\mathfrak{m}}}(k^n) \neq 0$, hence is a generator (it is finitely generated and projective by hypothesis of being Azumaya). This does not necessarily imply that $A \cong \mathrm{Mat}_n(R)$ as in the field case.

18.2 Cohomological description of the Brauer group over a ring - preliminary discussion

Recall that for a central simple algebra A over a field F, we proved that

a) A is split over some algebraic extension of F.

b) We can choose such an extension to be separable.

c) For a fixed splitting Galois extension E/F, the action of $G = \mathrm{Gal}(E/F)$ on $A \otimes_F E \cong \mathrm{Mat}_n(E)$ leads to the cohomological description of the Brauer group.

These statements can be generalized to a Noetherian commutative ring R.

18.3 Faithfully flat ring homomorphisms and faithfully flat descent

Let A be an Azumaya algebra over a Noetherian commutative ring R. The obvious analog of point 1 is some S with a map $R \to S$ such that A_S splits. However, in the ring setting, we can lose information from the base change map; for example, $R = S_1 \times S_2 \to S = S_1$. So we need an additional condition on S.

One that works well is that S is **faithfully flat** over R.

Definition 18.8: A ring S is **flat** over R if the functor $M \mapsto M_S$, R-Mod $\to S$-Mod is exact. It is **faithfully flat** over R if it is conservative, i.e., $M \to N$ is an isomorphism iff $M_S \to N_S$ is an isomorphism.

Remark 18.9: If S is flat, the conservativity condition is equivalent to $M_S \neq 0$ for $M \neq 0$.

Definition 18.10: Let M be an R-module. The **descent data** for M is a module $N = M_S$ over S, and an isomorphism ι between the base changes of N to $S \otimes_R S$ such that the three base changes to $S \otimes_R S \otimes_R S$ form a commutative diagram.

Proposition 18.11: If $R \to S$ is faithfully flat, the functor sending M to its descent data is an equivalence.

Remark 18.12: There is a parallel algebraic geometry statement. Notice that $S_1 \otimes_R S_2$ is the coproduct in the category of commutative rings, and Aff $=$ Comm$^{\mathrm{op}}$ (affine schemes) where R corresponds to $\mathrm{Spec}(R)$, so $\mathrm{Spec}(S_1 \otimes_R S_2) = \mathrm{Spec}(S_1) \times_{\mathrm{Spec}(R)} \mathrm{Spec}(S_2)$, the fiber product.
The descent data is parallel to how to define a vector bundle or sheaf on X by gluing the corresponding data for an open covering $X = \bigcup U_i$. Replace $\mathrm{Spec}(R)$ by X, $\mathrm{Spec}(S)$ by $U = \bigsqcup U_i$, and $S \otimes_R S$ is replaced by the disjoint union of $U_i \cap U_j$. The compatibility condition for base changes to $S \otimes_R S \otimes_R S$ correspond to checking that the data for $U_i \cap U_j \cap U_k$ makes sense.

Noetherian rings have a faithful flatness criterion.

Definition 18.13: A homomorphism of commutative rings $R \to S$ is **formally smooth** if for every commutative ring $T = \tilde{T}/I$ with $I^2 = 0$ and compatible maps $R \to \tilde{T}$, $S \to T$, you can lift to $S \to \tilde{T}$.

Example 18.14: Suppose R, S are finitely generated k-algebras with k a field, $T = k$, $\tilde{T} = k[t]/t^2$. Then $I = (t)$. We have maps $R \to k[t]/t^2$ and $S \to k$, and so also a map $R \to k$, so we have maximal ideals $\mathfrak{m}_S \subset S, \mathfrak{m}_R \subset R$ with residue field k. Extending a map $R \to k$ to a homomorphism to $k[t]/t^2$ is equivalent to specifying a vector in $(\mathfrak{m}_R/\mathfrak{m}_R^2)^*$, i.e. the **tangent vector** to $\mathrm{Spec}(R)$ at the corresponding point.
So formal smoothness implies that the map on tangent spaces induced by $R \to S$ is onto, a condition appearing in the definition of submersion in differential geometry.

Now suppose that R is Noetherian and S is finitely generated over S. If S is formally smooth over R, then it is flat over R. Moreover, if the map on k-points $\mathrm{Hom}(S, k) \to \mathrm{Hom}(R, k)$ is onto for every algebraically closed field k, then S is faithfully flat.

18.4 Universal splitting

Let A be an Azumaya algebra over R of constant rank d^2. Then we can construct an example of a faithfully flat ring S splitting R.

Theorem 18.15:

a) Consider the functor F sending a commutative R-ring S to the set of isomorphisms $A_S \cong \mathrm{Mat}_d(S)$. This functor is representable and is represented by a ring S_{univ} finitely generated over R.

b) S_{univ} is formally smooth over R. If R is Noetherian, it is faithfully flat over R.

Proof. Recall that A is a projective module over R (of rank d^2). So $A \cong e(R^N)$ as an R-module where $e \in \mathrm{Mat}_N(R)$ is an idempotent.
First consider the functor sending S to the S-module isomorphisms $A_S \cong S^{d^2}$. If $A \cong e(R^N)$; then such an isomorphism is equivalent to producing two matrices $i \in \mathrm{Mat}_{d^2,N}(S), j \in \mathrm{Mat}_{N,d^2}(S)$ such that $ij = I_{d^2}$ and $ji = e$. These are degree 2 equations in the entries of i, j, while the requirement that the isomorphism is compatible with the algebra structure on $\mathrm{Mat}_d(S)$ is another collection of degree 2 equations on the matrix entries of i. So we can define S_{univ} as the quotient of the polynomial ring in $2Nd^2$ variables over S by the ideal generated by these degree 2 equations.
To check that S_{univ} is formally smooth over R, we show that if $A_T \cong \mathrm{Mat}_n(T)$, then $A_{\tilde{T}} \cong \mathrm{Mat}_n(\tilde{T})$ where $T = \tilde{T}/I, I^2 = 0$, since that's what it means to be able to lift to a map $S_{\mathrm{univ}} \to \tilde{T}$. Consider a rank 1 idempotent $e \in \mathrm{Mat}_n(T)$ (without losing the generality we can assume that $e = e_{11}$). We will use the same notation for the corresponding element on A_T. So A_T maps isomorphically to $\mathrm{End}_T(A_T e)$. We can lift e to $\tilde{e} \in A_{\tilde{T}}$ such that $\tilde{e} \bmod I = e$. Then

$$A_{\tilde{T}} \to \mathrm{End}_{\tilde{T}}(A_{\tilde{T}} e)$$

is a map of free $\tilde{T}$-modules of rank d^2 that is an isomorphism modulo I, hence an isomorphism.
Hence, S_{univ} is flat over R when R is Noetherian. To check it is faithful, we need to check that $A_k \cong \mathrm{Mat}_d(k)$ for every algebraically closed field k, but this is one of the properties of Azumaya algebras. □

18.5 Rewriting cochain complex for $H^*(G, E^\times)$

Let's rewrite the complex used to compute $H^*(G, E^\times)$ for a finite Galois field extension E/F, $G = \mathrm{Gal}(E/F)$ in a way that can be generalized to Noetherian commutative rings. Recall that the nth term is $C^n = \mathrm{Map}(G^n, E^\times) = (\prod_{G^n} E)^\times$. From Galois theory, $\prod_G E \cong E \otimes_F E$ (this is an isomorphism of algebras). By induction, $\prod_{G^n} E = E \otimes_F E \otimes_F \cdots \otimes_F E$ where there are $n+1$ factors in the RHS. Thus

$$C^n = (E \otimes_F \cdots \otimes_F E)^\times$$

where there are $n+1$ factors in the RHS.

19 April 20 - Brauer group of a ring cont., localization

19.1 Amitsur cohomology

Let $F\colon \mathrm{Comm} \to \mathrm{Ab}$ be a functor. We can generalize the complex from the previous lecture to F, though we will mostly use $R \mapsto \mathbb{G}_m(R) = R^\times$. Given a homomorphism $R \to S$ we can form the **Amitsur complex** as follows:

Write $S_R^{\otimes n} = S \otimes_R \cdots \otimes_R S$ with n factors in the RHS. Set

$$C^n := F(S_R^{\otimes n+1}),\ d_n := \sum_{k=0}^{n+1} (-1)^k F(i_k)\colon C^n \to C^{n+1}$$

where $i_k\colon S_R^{\otimes n+1} \to S_R^{\otimes n+2}$ is the insertion map that puts a 1 in the kth place, i.e.

$$s_0 \otimes \cdots \otimes s_n \mapsto s_0 \otimes \cdots \otimes s_{k-1} \otimes 1 \otimes \cdots \otimes s_n.$$

We denote its cohomology by $H^i_{S/R}(F)$.

Example 19.1: Let $R = F, S = E$ with E/F a finite Galois extension and let the functor be $\mathbb{G}_m$. Recall that there is an isomorphism

$$(E^{\otimes_F n+1})^\times \xrightarrow{\sim} \left(\prod_{G^n} E\right)^\times = \mathrm{Map}(G^n, E^\times).$$

Choosing the isomorphism amounts to defining pairwise distinct homomorphisms

$$h_{g_1\cdots g_n}(x_0 \otimes x_1 \otimes \cdots \otimes x_n) = x_0 g_1(x_1) g_1 g_2(x_2) \cdots g_1 \cdots g_n(x_n)$$

where $h_{g_1\cdots g_n} : E_F^{\otimes n+1} \to E$. This commutes with i_k since if you let $x_k = 1$, you skip the $(k+1)$th factor and you get

$$h_{g_1\cdots g_n}(x_0 \otimes x_1 \otimes \cdots \otimes x_{k-1} \otimes 1 \otimes x_{k+1} \otimes \cdots \otimes x_n) = h_{g_1,\ldots,g_{i-2},g_{i-1}g_i,g_{i+1},\ldots,g_n}(x_0 \otimes \cdots \otimes x_n).$$

Hence the Amitsur complex is the standard complex computing $H^*(G, E^\times)$.

Remark 19.2: The algebraic geometry interpretation: Since $\mathrm{Comm}^{\mathrm{op}} = \mathrm{Aff}$, F can also be interpreted as a contravariant functor $\mathrm{Aff} \to \mathrm{Ab}$. Then $S \to R$ corresponds to $\mathrm{Spec}(R) \to \mathrm{Spec}(S)$; consider the analogous construction where you replace an affine scheme by a topological space, so we can instead consider morphisms $U \to X$ where U is the disjoint union $\bigsqcup U_i$ of open subsets in an affine covering of X. If the assignment of an abelian group to each $U_i \to X$ comes from a sheaf $\mathcal{F}$ on X, we recover the Cech complex for $H^*(X, \mathcal{F})$.

19.2 Relationship between Brauer group and Amitsur cohomology

We sketch how to correspond Azumaya algebras with a class in the second cohomology. Let A be an Azumaya algebra over R and choose an isomorphism $A_S \cong \mathrm{Mat}_n(S)$. Then we have two isomorphisms $A_{S_R^{\otimes 2}} \cong \mathrm{Mat}_n(S_R^{\otimes 2})$, and again, their ratio will be an Amitsur 1-cocycle c with nonabelian coefficients that is independent of the choice of isomorphism up to scaling. Hence it gives an element in $H^1_{S/R}(\mathrm{PGL}_n)$, where PGL_n is the functor $R \mapsto \mathrm{PGL}_n(R)$ (again, these will be nonabelian groups). Notice that $\mathrm{PGL}_n(R) = \mathrm{Aut}(\mathrm{Mat}_n(R))$ is an algebraic group and the homomorphism $\mathrm{GL}_n(R)/R^\times \to \mathrm{PGL}_n(R)$ may not be surjective (unlike in the field case).

Let's just assume that we can lift c to GL_n, e.g. the map $\mathrm{GL}_n(S \otimes_R S) \to \mathrm{PGL}_n(S \otimes_R S)$ is surjective, so c lifts to $\tilde{c} \in \mathrm{GL}_n(S \otimes_R S)$. Then we can get a cocycle in $H^2_{S/R}(\mathbb{G}_m)$ by the same procedure as in the field case: consider the differential of $\tilde{c}$, which takes values in $E^\times$, giving the desired cocycle.

Remark 19.3: In fact, one can find a faithfully flat S for which a lift $\tilde{c}$ exists, but the proof is beyond the scope of the lecture. Then you can define $H^i_{\mathrm{fl},A}(R, \mathbb{G}_m)$ (A for Amitsur) as $\mathrm{colim}_S H^i_{S/R}$ where the colimit is over all faithfully flat S. Restricting to étale S, you get $H^i_{\text{ét},A}(R, \mathbb{G}_m)$, and this coincides with the étale cohomology of $\mathrm{Spec}(R)$.
We have injective maps from $\mathrm{Br}(R)$ into $H^2_{\mathrm{fl},A}(R, \mathbb{G}_m)$ and $H^2_{\text{ét},A}(R, \mathbb{G}_m)$.

19.3 Final remarks on Brauer group

First, we describe how to generalize separable splittings to rings. It turns out that for an Azumaya algebra A over R, we can always find an étale, faithfully flat homomorphism $R \to S$ such that A_S splits.

Definition 19.4: A ring homomorphism $R \to S$ is **étale** if for every commutative ring $T = \tilde{T}/I$ with $I^2 = 0$ and compatible maps $R \to \tilde{T}$, $S \to T$, there exists a unique compatible map $S \to \tilde{T}$.

Exercise 19.5: A finite field extension is étale iff it is separable.

Theorem 19.6: Let R be a (formally) smooth finitely generated commutative domain over an algebraically closed field and $F = \mathrm{Frac}(R)$. Then $\mathrm{Br}(R) \hookrightarrow \mathrm{Br}(F)$.

Proof (Sketch). The proof involves an object called the **Brauer-Severi variety** (to be denoted by B). We need the notion of a line bundle (a locally free coherent sheaf of rank 1) and the fact that for a smooth variety X over a field and $U \subset X$ an open subvariety, every line bundle on U can be extended to one on X. This follows from the correspondence between line bundles and divisors and the fact that the closure of a divisor on U is a divisor on X. We also need the concept of an algebraic group action on an algebraic variety and the quotient by such an action. Let A be an Azumaya algebra on $X = \operatorname{Spec}(R)$ and $S = S_{\text{univ}}$ be the universal splitting ring. Then $G = \mathrm{PGL}_n$ acts on $Y = \operatorname{Spec}(S)$ so that $Y/G \cong X$. Recall that G also acts on $\mathbb{P}^{n-1}$. Set

$$B := (\mathbb{P}^{n-1} \times Y)/G.$$

Thus $B \to X$ and every geometric fiber of this map is isomorphic to $\mathbb{P}^{n-1}$. Then one can check that A is split iff there exists a line bundle L on B whose restriction to a geometric fiber is isomorphic to the line bundle $O(1)$ on $\mathbb{P}^{n-1}$. If A_F splits then there exists a nonempty open $U \subset X$ such that A_U splits, so A splits. □

19.4 Localization

Let R be a ring and S a multiplicatively closed subset, i.e. $1 \in S$ and $a, b \in S \Rightarrow ab \in S$.

Definition 19.7: The **localization** R_S of R at S is the universal ring receiving a homomorphism from R sending S to invertible elements. That is,

$$\operatorname{Hom}(R_S, T) = \{f : R \to T \mid f(s) \text{ is invertible } \forall s \in S\}.$$

The Yoneda Lemma shows that R_S is unique up to unique isomorphism if it exists.

Lemma 19.8: $R_S = R\langle t_s \rangle_{s \in S}/(t_s s = s t_s = 1)$.

19.5 Ore conditions

Unlike in the commutative ring case, it is hard to say much about R_S from this construction; for example, we don't even know if R_S is the zero ring. We can impose additional conditions on S to give R_S an explicit description.

Definition 19.9: Let $S \subset R$ be a multiplicative subset. The **(right) Ore conditions** are

- (O1) For all $a \in R, s \in S$, then $aS \cap sR \neq \emptyset$.
- (O2) For all $a \in R, s \in S$, if $sa = 0$, then there exists $t \in S$ such that $at = 0$.

If S satisfies O1, it is called a **right Ore set**. If S satisfies O1 and O2, it is called a **right reversible** or **right denominator set**. There are analogous definitions for left everything.

Remark 19.10: O1 allows us to pull denominators of fractions to the right: if $aS \cap sR \neq \emptyset$, then $at = sb$ for $t \in S, b \in R$. So using formal inverses, $s^{-1}a = bt^{-1}$.

Using O1 and O2, then R_S will consist of pairs $(a, s) \in R \times S$ modulo the equivalence that $(a, s) \sim (a', s')$ if there exist $u, u' \in R$ such that

$$au = a'u', \; su = s'u' \in S.$$

That is,

$$as^{-1} = (au)(su)^{-1} = (a'u')(s'u')^{-1} = a'(s')^{-1}.$$

This has a ring structure where $a \mapsto (a, 1)$ is a ring homomorphism.

Remark 19.11: Localization of a ring or a module can also be presented as a filtered colimit. We can create a diagram category D where the objects are S and $\operatorname{Hom}(s, t) = \{u \mid su = t\}$ and composition is given by $v \circ u = uv$. Then if O1 and O2 both hold, then D is filtered. Moreover, R_S is the filtered colimit $\lim_D R$. This shows that localization is exact because filtered colimits are (for abelian groups); also, it comes with the forgetful functor. We will prove this next lecture.

20.1 Ore localization and regular elements

Proposition 20.1: Let S be a right reversible multiplicative subset in a ring R, i.e. it satisfies O1 and O2. Say that $(a, s) \sim (a', s')$ if there exist $t, t' \in R$ such that $at = a't'$ and $st = s't' \in S$ (that is, $a/s = a'/s'$). This is an equivalence relation on $R \times S$ and the map $(a, s) \mapsto as^{-1}$ is a bijection between $(R \times S)/\sim$ and the localization R_S.

Proof. The relation is clearly reflexive and symmetric, we need to show transitivity. Suppose $(a, s) \sim (a', s')$, so $at = a't'$ and $st = s't' \in S$ for some $t, t' \in R$, and also $(a', s') \sim (a'', s'')$, so there exist $u, u' \in R$ such that $a''u = a'u', s''u = s'u' \in S$. We need to find $v, v'' \in S$ such that $av = a''v'', sv = s''v'' \in S$.
Apply O1 to $\alpha := s't', \sigma := s'u'$ to see that there exists $z_0 \in S, x_0 \in R$ such that $s't'z_0 = s'u'x_0$. Applying O2 to $s'(t'z_0 - u'x_0) = 0$, there exists some $r \in S$ such that $(t'z_0 - u'x_0)r = 0$. In other words, there exist elements $z \in S$, $x \in R$ satisfying $t'z = u'x$.
Therefore,

$$atz = a't'z = a'u'x = a''ux$$

with

$$stz - s''ux = s'(t'z - u'x) = 0 \Rightarrow s''ux = stz \in S.$$

Hence, $\sim$ is an equivalence relation.
To define a ring structure on the set of equivalence classes, write as^{-1} for the equivalence class of (a, s). To multiply $as^{-1} \cdot bt^{-1}$, find $c \in R, u \in S$ with $bu = sc$ and set

$$as^{-1} \cdot bt^{-1} = ac(tu)^{-1}.$$

To add $as^{-1} + bt^{-1}$, find s', t' such that $ss' = tt' \in S$ (these exist using O1), then

$$as^{-1} + bt^{-1} = (as')(ss')^{-1} + (bt')(tt')^{-1} = (as' + bt')(ss')^{-1}.$$

One can check that these are well-defined and produce an associative ring. Denote this ring by RS^{-1}. There is a map $R_S \to RS^{-1}$ since the map $R \to RS^{-1}$ sending $r \mapsto (r, 1)$ sends S to units. In the other direction, there is a map $RS^{-1} \to R_S$ sending $(a, s) \mapsto as^{-1}$. It's easy to see this map is a homomorphism and the two homomorphisms above are inverse isomorphisms. □

Corollary 20.2: For a right denominator set $S \subset R$, the kernel of the canonical homomorphism $R \to R_S$ is the set of elements whose right annihilator intersects S.

Proof. The kernel is the set of elements $a \in R$ such that $(a, 1) \sim (0, 1)$, which is true iff $as = 0$ for some $s \in S$. □

Definition 20.3: An element of R is **regular** if it is neither a left nor right zero divisor.

Corollary 20.4: If S consists of regular elements, the natural map $R \to R_S$ is injective.

20.2 Ore localization as a filtered colimit

Extending the remark 19.11 from last time, the localization can also be interpreted as a filtered colimit.

Recall from Definition 9.12 that a category D is filtered if $\mathrm{Ob}(D) \neq \emptyset$ and

- for every $a, b \in D$, there exists $c \in D$ such that $\mathrm{Hom}(a, c)$ and $\mathrm{Hom}(b, c)$ are nonempty
- for every pair of parallel morphisms $e, f \colon a \to b$, there exists $g \colon b \to c$ such that $g \circ e = g \circ f$.

Taking the filtered limit of abelian groups is exact and commutes with the filtered colimit of sets under the forgetful functor. The filtered colimit of sets can be described as follows: for a functor $F\colon D \to \mathrm{Set}$, its colimit is the quotient

$$\bigsqcup_{a \in \mathrm{Ob}(D)} F(a)/{\sim}$$

where $x \sim y$ for $x \in F(a), y \in F(b)$ if $y = F(e)(x)$ for some $e \in \mathrm{Hom}(a, b)$. (That is, there's an arrow in the image of F from x to y.)

As in the last lecture, we can create a diagram category D where the objects are S and $\mathrm{Hom}(s, t) = \{u \mid su = t\}$ and composition is given by $v \circ u = uv$.

Proposition 20.5: If S is a right denominator set (i.e. both O1 and O2 hold), then D is filtered.

Proof. First, D is nonempty because $1 \in S$.
For every $s, t \in \mathrm{Ob}(D) = S$, then by O1 we have $S \cap t^{-1}sR \neq \emptyset$, hence there exists a, b such that $sa = tb$, so $\mathrm{Hom}(s, sa)$ and $\mathrm{Hom}(t, tb)$ are nonempty.
Two parallel morphisms $s \to t$ are $a, b \in R$ such that $t = sa = sb$. Then by O2, $s(a - b) = 0$ implies there exists $u \in S$ such that $(a - b)u = 0$. So by composing the two parallel morphisms a, b with the morphism $t \to tu$ given by u, we get the same morphism. □

Now for M a right R-module, define a functor $F_M\colon D \to R^{\mathrm{op}}\text{-mod}$ by sending every object to M and every morphism corresponding to $u \in R$ to right multiplication by u. Hence, R_S is the colimit of F_R. Therefore,

$$\operatorname{colim} F_M := M \otimes_R R_S =: M_S$$

is the localization of M at S, and $M \mapsto M_S$ is exact.

Example 20.6: Let $R = k[x]$ and $S = \{1, x, x^2, \dots\} \subset R$ be the powers of x. This gives the filtered category D whose objects are the elements of S, and there's a map $x^i \to x^j$ iff $i \leq j$; this map is precisely given by multiplication by x^{j-i}. Then we can form the localization R_S by taking the colimit over the functor F_R: indeed,

$$k[x, x^{-1}] = k[x]_x R_S = \operatorname{colim}\left(k[x] \xrightarrow{\cdot x} k[x] \xrightarrow{\cdot x} k[x] \xrightarrow{\cdot x} \cdots\right).$$

This gives us the sequence of injections

$$k[x] \hookrightarrow k\{x^{-1}, 1, x, \dots\} \hookrightarrow k\{x^{-2}, x^{-1}, 1, x, \dots\} \hookrightarrow k\{x^{-3}, x^{-2}, \dots\} \hookrightarrow \dots$$

for which the colimit (we can interpret this as basically the union) is indeed $k\{x^n \mid n \in \mathbb{Z}\}k[x, x^{-1}] = R_S$.

Remark 20.7: Ore conditions can also be generalized to categories: many important constructions involve inverting a class of morphisms in a category, and the generalization of the Ore conditions guarantees a manageable result. The construction of a **derived category** as a localization of the homotopy category of complexes is an example.

20.3 Ore domains

Definition 20.8: A ring R is an **Ore domain** if it's a domain and $R \backslash \{0\}$ satisfies O1. In this case, R_S for $S = R \backslash \{0\}$ is clearly a skew field and $R_S = \mathrm{Frac}(R)$.

Example 20.9: A free ring (e.g. over a field) with at least two generators is *not* an Ore domain: if x, y are free generators then $xR \cap yR = 0$.

Proposition 20.10: Assume R is a domain.

a) (Goldie) Either R is a right Ore domain or it contains a free right ideal of infinite rank.

b) (Jategoankar) Say R is an algebra over a field k. Then either R is a left and right Ore domain or it contains a free ring $k\langle x, y\rangle$.

Proof. *a*) Suppose R is not a right Ore domain, so there exist a, b such that $aS \cap bR = \emptyset$ (recall that $S = S \setminus \{0\}$). Then we claim that $a, ba, b^2a, \ldots,$ is right independent over R. Otherwise, we could find $\{r_i\}$ such that

$$\sum_{i=0}^{n} b^i a r_i = 0 \Longrightarrow -ar_0 = b\left(\sum_{i=1}^{n} b^{i-1} a r_i\right),$$

contradiction (note that we can assume that $r_0 \neq 0$ i.e. $-r_0 \in S$).

b) Suppose R is not a right Ore domain and pick x, y such that $xR \cap yR = 0$. Let $f(x, y) = a + xf_1 + yf_2$ be a minimal relation where $a \in k$. If $a = 0$, then $xf_1 = yf_2 \neq 0$ but $xR \cap yR = 0$, contradiction. If $a \neq 0$, multiplying everything by y on the right, we have $ay + xf_1y + yf_2y = 0$. Since $a \in k$, $ay = ya$ and $x(f_1y) = y(a + f_2y)$. These are again both nonzero: if $f_1y = 0$, then $f_1 = 0$ because R is a domain, so $yf_2 + a = 0$, so y is invertible. Then $yR = R$, so $xR \cap yR \neq 0$, contradiction. Likewise, $a + f_2y \neq 0$. So $xR \cap yR$ has a nonzero element, a contradiction. Thus x, y generate a free algebra.
The same argument works if R is not a left Ore domain.
On the other hand, if R is a right Ore domain, then any x, y satisfying $k\langle x, y\rangle \subseteq R$ must have $xR \cap yR = 0$, which contradicts the Ore assumption.
In conclusion, either R is a right and left Ore domain, OR R contains a subalgebra of the form $k\langle x, y\rangle$.

□

20.4 Growth of algebras

Let A be a finitely generated k-algebra for a field k. Let V be a (finite-dimensional) vector space of generators for A, so we have an onto map $TV \twoheadrightarrow A$ where TV is a tensor algebra. Let $A^V_{\leqslant n}$ be the image of $\bigoplus_{i \leqslant n} V^{\otimes i}$ and set

$$d_V(n) := \dim_k(A^V_{\leqslant n}).$$

For a different space of generators W, $d_W \neq d_V$, but $d_W(n) \leqslant d_V(n_0 n)$ always for some fixed n_0 because $A^W_{\leqslant n} \subset A^V_{\leqslant n_0 n}$ for some n_0.

So say that two (monotone) functions f, g on $\mathbb{N}$ are equivalent if there exists n_0 such that

$$f(n) \leqslant g(n_0 m),\ g(n) \leqslant f(n_0 m).$$

So the equivalence class of $d_V(n)$ is independent of the choice of V.

Definition 20.11: We say that A has **exponential growth** if $d(n) \geqslant c\alpha^n$ for some constants $\alpha > 1, c$. If A does not have exponential growth, it necessarily has **subexponential growth**, i.e. for all $\alpha > 1$, $f(n)/\alpha^n \to 0$.

Example 20.12: If A contains a free algebra, then A has exponential growth.

Corollary 20.13 (of Proposition 20.10): If A is a domain of subexponential growth, then A is an Ore domain.

Example 20.14: The Weyl algebra

$$W_n = k\langle x_1, \ldots, x_n, y_1, \ldots, y_n\rangle / ([x_i, y_j] = \delta_{ij}, [x_i, x_j] = [y_i, y_j] = 0)$$

and $U(\mathfrak{g})$ for $\mathfrak{g}$ a finite-dimensional Lie algebra are domains of polynomial, hence subexponential, growth, and therefore are Ore domains.

20.5 Semi-prime rings and Goldie's theorem

Recall that an element is regular if it's neither a left or right zero divisor.

Remark 20.15: For a regular element, left invertibility is equivalent to right invertibility, since $sr = 1 \Rightarrow rsr = r \Rightarrow rs = 1$.

Definition 20.16: A ring is called **prime** if $IJ \neq 0$ for any two nonzero two-sided ideals $I, J \subset R$. It is **semi-prime** if $I^2 \neq 0$ for any nonzero two-sided ideal $I \subset R$.

Recall that a ring is semi-primitive if its Jacobson radical vanishes, which is equivalent to the existence of a faithful semisimple (either left or right) module.

Proposition 20.17: Every semi-primitive ring is semi-prime.

Proof. Suppose $I \subset R$ is a nonzero two-sided ideal, and R is semi-primitive. So we can find an irreducible R-module L such that $IL \neq 0$. Then from the density theorem, it follows that we can find $x \in I, v \in L, v \neq 0$ where $xv = v$. Hence, $x^2 \neq 0$. □

The converse is not true, but we do have the following:

Theorem 20.18 (Goldie): If R is a semi-prime right Noetherian ring, then the set S of all regular elements satisfies (right) O1, and $Q = R_S$ is an Artinian semisimple ring.

Corollary 20.19: If R is left or right Noetherian, it admits a homomorphism to $\mathrm{Mat}_n(D)$, so it satisfies the IBN.

Proof. If R is right Noetherian, then $\bar{R} := R/J(R)$ is semi-primitive and right Noetherian, hence semi-prime. By Goldie's theorem, $\bar{R}_S$ is Artinian semisimple, so $\bar{R}_S = \prod_{i=1}^n \mathrm{Mat}_{d_i}(D_i)$. Hence

$$R \longrightarrow \bar{R} \longrightarrow \bar{R}_S \longrightarrow \mathrm{Mat}_{d_1}(D_1)$$

is the desired homomorphism. □

The idea of the proof of the theorem is that sR is "too big" to miss aS; we need a notion of size.

Definition 20.20: Let M be a right R-module. A submodule $E \subset M$ is **essential** if for all nonzero $N \subset M$, $N \cap E \neq 0$. That is, every nonzero submodule in M has a nonzero intersection with N. We say that M is **uniform** if $M \neq 0$ and every nonzero submodule in M is essential.

Example 20.21: If M is of finite length, $E \subset M$ is essential iff $E \supset \mathrm{Soc}(M)$ and M is uniform iff $\mathrm{Soc}(M)$ is simple. For example, for $R = k[t]$, $M = k[t]/(t^n)$ is uniform. Another example is a domain R considered as a (right) module over itself.

Lemma 20.22: If $N \subset M$ is a submodule, then there exists a submodule $N' \subset M$ such that $N \oplus N'$ is an essential submodule in M. Then N' is called the **essential complement** of N.

Proof. Consider all submodules with zero intersection with N. Then the condition of Zorn's Lemma holds, so there exists a maximal element N' in this set. Then $N \oplus N'$ is essential in M. □

The measure of size we will use is the maximal number of uniform submodules of M such that their direct sum is also a submodule of M.

Proposition 20.23:
a) Let M be a Noetherian module. Then it contains an essential submodule that is a sum of uniform submodules, $E = \bigoplus_{i=1}^{n} U_i$, E essential and U_i uniform.
b) The number of uniform summands is independent of choices and is the **Goldie rank** or **uniform dimension**.
c) Every submodule of full Goldie rank is essential. That is, if $M \supset N$ and $\mathrm{Grank}(M) = \mathrm{Grank}(N)$, then N is essential in M.

Corollary 20.24: If $s \in R$ is a regular element, then $sR \subset R$ is an essential ideal.

Lemma 20.25: The preimage of an essential submodule is essential.

Proposition 20.26: An essential right ideal in a semi-prime right Noetherian ring contains a regular element.

Next time, we will prove these and discuss other facts about essential modules.

21.1 More on essential modules

Corollary 21.1: A module M has no proper essential submodules iff it is semisimple.

Proof. We proved that a module M is semisimple iff every submodule N has a direct complement. So if $N \subset M$, we know it has an essential complement N' such that $N \oplus N'$ is essential. If M has no proper essential submodules, then $N \oplus N' = M$ and M is semisimple. If M is semisimple, every submodule's direct complement doesn't intersect it, so there are no proper essential submodules. □

Lemma 21.2:
a) If $M \supset N \supset P$ with N essential in M and P essential in N, then P is essential in M.
b) The preimage of an essential submodule is essential.
c) If $N_1 \subset M_1, N_2 \subset M_2$ are essential, then $N_1 \oplus N_2 \subset M_1 \oplus M_2$ is essential.

Proof. *a)* If $S \subset M$ has nonzero intersection with N (use that $N \subset M$ is essential), then $S \cap N \subset N$ has nonzero intersection with P (use that $P \subset N$ is essential).
b) Let $\varphi \colon M \to N$ and $E \subset N$ essential. Suppose $V \subset M$ is a nonzero submodule. Then either $V \subset \ker(\varphi)$ or $\varphi(V)$ is nonzero. If $V \subset \ker(\varphi)$, $V \subset \varphi^{-1}(E)$. If $\varphi(V) \neq 0$, then $\varphi(V) \cap E \neq 0$, so $V \cap \varphi^{-1}(E) \neq 0$.
c) By a), it's enough to consider $M_1 = N_1$. Then $M_1 \oplus N_2$ is the preimage of N_2 under the projection $M_1 \oplus M_2 \twoheadrightarrow M_2$, so it is essential by b).

□

21.2 Goldie rank

Definition 21.3: A module M has **finite Goldie rank** if it does not contain an infinite (direct) sum of nonzero submodules.

Example 21.4: If M is Noetherian, it has a finite Goldie rank. In fact, one can restate the finite Goldie rank condition as the condition that split increasing chains of submodules should stabilize, where a chain of submodules M_i splits if M_i has a direct complement in M_{i+1} for all i.

Proposition 21.5: A finite Goldie rank module contains an essential submodule which is a finite (direct) sum of uniform submodules.

Proof. Suppose for contradiction that M does not contain such an essential submodule. Then M is not uniform (as otherwise M itself would be uniform, hence an essential submodule which is trivially a one-term direct sum of a uniform submodule), so it has a nonessential submodule N_1 with essential complement C_1. Since we assumed that M does not contain any essential submodules which are a finite direct sum of uniform submodules, it follows that N_1 and C_1 cannot both be uniform, else $N_1 \oplus C_1$ would give such an essential submodule. So without loss of generality suppose C_1 is not uniform. Then repeat the same argument for C_1; we get two submodules N_2, C_2 where $N_2 \oplus C_2 \subset C_1$ and at least one of N_2, C_2 are not uniform (without loss of generality, say C_2 is not uniform). Thus by induction we get $N_1, N_2, \ldots$ where $M \supset N_1 \oplus N_2 \oplus \cdots$, contradicting the assumption that M has finite Goldie rank. □

Theorem 21.6: Suppose M has finite Goldie rank and contains an essential submodule $E = \bigoplus_{i=1}^{m} U_i$ which is a finite direct sum of uniform submodules. If $M \supset N = \bigoplus_{i=1}^{n} N_i$ with $N_i \neq 0$, then $n \leqslant m$. If $m = n$, then N is essential and each N_i is uniform. In this case, we say that M has **Goldie rank** m, abbreviated by Grank.

Proof. First, $N' := \bigoplus_{i=2}^{n} N_i$ is not essential, since it doesn't meet N_1. Then we claim that $N' \cap U_i = 0$ for some i. Otherwise, $N' \cap U_i \neq 0$ is essential in U_i, so by the lemma $\bigoplus_{i=1}^{m}(N' \cap U_i)$ is essential in E, hence essential in M, and therefore N' is essential in M.
WLOG say that $N' \cap U_1 = 0$. Then $U_1 \oplus N_2 \oplus \cdots \oplus N_n \subset M$. Suppose that $n > m$. Then continuing inductively, with possible reindexing, $U_1 \oplus \cdots \oplus U_m \oplus N_{m+1} \oplus \cdots \oplus N_n = E \oplus N_{m+1} \oplus \cdots \oplus N_n \subset M$, contradicting that E is essential. Therefore, $n \leqslant m$.
If $m = n$, then N is essential. If not, we'd have an essential complement S and $N \oplus S$ would be an essential sum of $n + 1$ nonzero submodules, contradiction. Likewise, each N_i is uniform: otherwise, it would have a nonessential submodule N_i' with essential complement N_i'', so we would again get a direct sum of $n + 1$ submodules. □

Corollary 21.7: If M has finite Goldie rank n, then every submodule in M with the same Goldie rank n is essential.

Corollary 21.8: The Goldie rank can also be defined as the maximal number of $M_i \neq 0 \subset M$ such that $\bigoplus_i M_i \subset M$.

Example 21.9: For semisimple modules, the Goldie rank is the number of simple summands.

21.3 Regular elements in essential ideals

Remark 21.10: Suppose $S \subset R$ consists only of regular elements. Then the localization of an essential (resp. uniform) ideal at S is essential (resp. uniform).

Theorem 21.11: An essential right ideal in a semi-prime, right Noetherian ring contains a regular element.

This will imply the first statement in Goldie's theorem 20.18: let S be the regular elements. Given $s \in S$, $sR \cong R$ so it has the same Goldie rank as R (as a right module over itself) and is essential in R (use Corollary 21.7). Hence, for any $a \in R$, the preimage of sR under the map $x \mapsto ax$ is an essential right ideal (Lemma 21.2) and contains a regular element t. Thus $aS \cap sR \neq \varnothing$, which implies O1; O2 is vacuous for regular elements.

To prove the theorem, we first start with a weaker claim.

Lemma 21.12: Let R be a right Noetherian, semi-prime ring and $I \subset R$ an essential right ideal. Then the left annihilator of I is zero.

Proof. Let J be the left annihilator of I. We know $J^2 \neq 0$ because R is semi-prime (if $I^2 = 0$, then $(JR)^2 = 0$ for the two-sided ideal JR). Replace I by $\mathrm{rAnn}(J)$; WLOG we can assume that I is maximal among right annihilators using the Noetherian property.
Since $J^2 \neq 0$, pick $x, y \in J$ such that $xy \neq 0$. Then $yR \cap I \neq 0$ since I is essential, so there exists r with $yr = z \in I$

and $xyr = 0$. Then

$$r \notin \mathrm{rAnn}(I), r \in \mathrm{rAnn}(xy) \Rightarrow \mathrm{rAnn}(xy) \supsetneq \mathrm{rAnn}(y) \supset I$$

which contradicts the maximality of I. □

Proposition 21.13: Any right ideal I contains an element x with $\mathrm{rAnn}(x) \cap I = 0$.

This proposition implies Theorem 21.11. Let I be an essential ideal. Then we can find $r \in I$ with $\mathrm{rAnn}(x) \cap I = 0$. Since I is essential, this means $\mathrm{rAnn}(r) = 0$ and rR is free. In particular, it has the same Goldie rank as R, so rR is essential in R. Then by the lemma, $\mathrm{lAnn}(rR) = \mathrm{lAnn}(r) = 0$. So r is regular.

22 May 2 - Goldie Theorem, PI rings

22.1 Finishing up Goldie Theorem

Proof (of Proposition 21.13). First, we prove the claim when I is uniform (see Definition 20.20). Again, $I^2 \neq 0$ since R is semi-prime, so pick $x, y \in I$, $xy \neq 0$. Then we claim that $\mathrm{rAnn}(x) \cap I = 0$. Otherwise, $\mathrm{rAnn}(x) \cap I$ is essential in I. Consider the homomorphism of (right) R-modules $L_y : R \to I$ given by $z \mapsto yz$. It follows from Lemma 21.2 that the preimage $L_y^{-1}(\mathrm{rAnn}(x) \cap I)$ is essential in R. So $\{z \in R \mid yz \in \mathrm{rAnn}(x)\}$ is essential in R. But then its left annihilator is zero by the above lemma, but $x \neq 0$ is in the annihilator, contradiction.
In general, choose a maximal subideal $J \subset I$ such that there exists $v \in J$ with $\mathrm{rAnn}(v) \cap J = 0$ (via the right Noetherian property). If $\mathrm{rAnn}(v) \cap I \neq 0$, pick a uniform ideal $U \subset \mathrm{rAnn}(v) \cap I$. There exists $u \in U$ with $\mathrm{rAnn}(u) \cap U = 0$. Set $x = u + v$.
Since $U \subset \mathrm{rAnn}(v)$, $U \cap J = 0$. So if $x \in \mathrm{rAnn}(u+v)$, then $x \in \mathrm{rAnn}(u) \cap \mathrm{rAnn}(v)$. Suppose $x = u' + v' \in U \oplus J$. Then $uu' + uv' = 0, vu' + vv' = 0$. But $vu' = 0$ since $U \subset \mathrm{rAnn}(v)$, so $vv' = 0 \Rightarrow v' = 0$. So $uu' = 0$ and $u' = 0$ by assumption on u. Thus, $J \oplus U$ is a larger subideal in I containing an element $u+v$ whose right annihilator has zero intersection with the ideal, contradicting the maximality of J. □

Proof (of Theorem 20.18). To finish proving the Goldie theorem, we need to show that R_S is Artinian semisimple. This is equivalent to R_S being semisimple as a right module over itself, which is equivalent to saying that R_S has no proper essential ideals. Suppose that $I \subset R_S$ is essential. Then $I \cap R$ is essential in R: $R \hookrightarrow R_S$ because S consists of regular elements, so the preimage of $I \subset R_S$, which is $I \cap R$ is essential.
Then $I \cap R$ contains a regular element (Theorem 21.11), i.e., $R \cap I \cap S$ is nonempty, so $I = R_S$. □

22.2 Goldie rings

The statement of Goldie's theorem required R to be semi-prime right Noetherian. However, the proof only uses the fact that R has 1) finite Goldie rank as a right module over itself (split ascending chains of right ideals stabilize) and 2) chains of right annihilators stabilize.

This is because even though we invoked the Noetherian property to find a maximal ideal $J \subset I$ with $v \in J$ such that $\mathrm{rAnn}(v) \cap J = 0$, the proof found an ideal of the form $J \oplus U$, so it suffices to use that split chains terminate.

Definition 22.1: If R has finite Goldie rank as a right module over itself and chains of right annihilators stabilize, we say that R is a **(right) Goldie ring**.

Example 22.2: Not every right Goldie ring is right Noetherian. For example, every commutative domain where every annihilator of a nonzero element is zero and every nonzero ideal is essential is a right Goldie ring but not necessarily right Noetherian.

22.3 Applications of Goldie's Theorem

Proposition 22.3: Let R be a semi-prime Goldie ring and S the set of its regular elements. Then if $I \subset J$ is an essential subideal, the localizations I_S and J_S coincide. Also, if I is uniform then I_S is irreducible.

Proof (Sketch). Essential embeddings and uniformity survive after localization. Over semi-simple Artinian rings, uniform modules are irreducible and essential embeddings are isomorphisms. □

Hence, Goldie rank is a measure of the size of an infinite-dimensional algebra (say, for algebras over a field) and it's an interesting question to understand it better and compare it with other measures.

Example 22.4: What is the Goldie rank of R as a module over itself? For example, if R is prime (in particular, if it is primitive), then $R_S \cong \mathrm{Mat}_n(D)$, and the Goldie rank will be n.

A very interesting story is related to the study of this invariant for $R = U(\mathfrak{g})/I$ where $\mathfrak{g}$ is a complex simple finite-dimensional Lie algebra (e.g. $\mathfrak{sl}(n)$) and I is a primitive ideal. Then the answer is given by the "Goldie rank polynomial"; the classification of ideals involves a parameter λ on which the answer depends polynomially. This is largely understood due to the work of various authors, including David Vogan, George Lusztig, Tony Joseph, and, more recently, Ivan Losev.

Another famous question related to noncommutative localization and Lie theory is the Gelfand-Kirillov conjecture. This states that for a large class of Lie algebras, including those mentioned above, the fraction field of $U(\mathfrak{g})$ (a domain of polynomial growth, hence an Ore domain) is isomorphic to the fraction field of a ring of the form $W_n[x_1, \ldots, x_r]$ where W_n is the Weyl algebra. This turned out to be false in general, but true for $\mathfrak{g} = \mathfrak{sl}(n)$. However, if $\bar{U} = U(\mathfrak{g})/\mathfrak{m}U(\mathfrak{g})$, where $\mathfrak{m}$ is a maximal ideal in the center of $U(\mathfrak{g})$, then the fraction field of $\bar{U}$ is indeed isomorphic to the fraction field of W_n for every simple complex Lie algebra.

22.4 PI rings

Definition 22.5: A ring R is a **polynomial identity (PI) ring** if there exists a nonzero element in the free algebra $P \in \mathbb{Z}\langle x_1, \ldots, x_n \rangle$ such that $P(r_1, \ldots, r_n) = 0$ for all $r_1, \ldots, r_n \in R$ (i.e., there is a polynomial identity that all elements satisfy).
Likewise, if A is an algebra over a field (or commutative ring) k, it is a **polynomial identity (PI) algebra** if there exists a nonzero $P \in k\langle x_1, \ldots, x_n \rangle$ such that any evaluation of P in A vanishes.

Example 22.6: Commutative rings are PI rings: take $P(x, y) = xy - yx$.

Example 22.7: Boolean rings (rings where every element is idempotent) are also PI rings with $P(x) = x^2 - x$.

Example 22.8: Let

$$S_n(x_1, \ldots, x_n) = \sum_{\sigma \in S_n} (-1)^{|\sigma|} x_{\sigma(1)} \cdots x_{\sigma(n)}.$$

We claim that this holds in every finite-dimensional algebra A over a field k of char $k \neq 2$ when $n > \dim_k(A)$. This is because evaluation of S_n is a skew-symmetric multilinear functional, hence is a map $\Lambda^n(A) \to A$. But if $n > \dim_k(A)$, then $\Lambda^n(A) = 0$.

22.5 Amitsur-Levitzki Theorem

Theorem 22.9 (Amitsur-Levitzki): The identity S_{2n} holds in the ring $\mathrm{Mat}_n(R)$ for any commutative ring R. Moreover, no (nonzero) homogeneous identity of smaller degree holds (assuming $R \neq 0$).

The second part of the theorem is easier and follows from the next two lemmas.

Lemma 22.10 (Staircase Lemma): $\mathrm{Mat}_n(R)$ does not satisfy a multilinear identity of degree $d < 2n$.

Proof. Consider the following $2n - 1$ elementary matrices:

$$e_{11}, e_{12}, e_{22}, \ldots, e_{n-1,n-1}, e_{n-1,n}, e_{n,n}$$

Their product in this order is an elementary matrix, namely e_{1n}, but their product in any other order vanishes. The first r matrices in that list for $r < 2n-1$ satisfy the same property.
A multilinear polynomial is a linear combination of multi-homogeneous monomials with coefficients in R. If a degree r monomial $x_1 \cdots x_r$ is in the polynomial, substitute the above elementary matrices for x_i and zero for the other variables (if any). Then our sum has exactly one nonzero summand, so the sum is nonzero. □

Lemma 22.11:

a) If a ring satisfies an identity P of degree d, then it satisfies a multilinear identity of the same degree.

b) If an algebra A over an infinite field k satisfies a polynomial identity $P = \sum P_d$ where P_d is homogeneous of degree d, then each P_d is also an identity satisfied by A.

Proof. *a)* Let $P = P(x_1, \ldots, x_n)$ be a degree d identity. We do double induction on the top degree of P in each variable and the number of variables in which it has that degree. Suppose $r > 1$ is the top degree and WLOG that P has degree r in x_1. Then consider

$$Q(x_0, \ldots, x_n) = P(x_0 + x_1, x_2, \ldots, x_n) - P(x_0, x_2, \ldots, x_n) - P(x_1, \ldots, x_n).$$

Q holds in our ring and has degree less than r in both x_0, x_1. For the other variables, their degree is most that of P. Note that Q is not identically zero: this is because for monomials M of degree d, the noncommutative polynomials

$$M' = M(x_0 + x_1, x_2, \ldots, x_n) - M(x_0, x_2, \ldots, x_n) - M(x_1, \ldots, x_n)$$

are linearly independent over R. This is because the monomials in M' which are linear in x_0 will enter M' with multiplicity 1, and we can reconstruct M from such a monomial by replacing x_0 by x_1.
Therefore, by induction we can find an identity P which has degree one in each variable. Suppose there is a variable x_i appearing in P in which P is not linear (so x_i appears in some monomials but not in others). Then

$$P(x_1, \ldots, x_i, \ldots, x_n) - P(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n)$$

is also an identity and is nonzero and linear in x_i. Repeating this inductively, we get a multilinear identity of the same total degree.

b) For $\lambda \in k$, $P_\lambda = \sum \lambda^d P^d$ is also a polynomial identity. Choosing distinct $\lambda_1, \ldots, \lambda_n$ with $n > \deg(P)$, the linear span of P_{λ_i} will contain P_d because the Vandermonde determinant doesn't vanish.

□

This furnishes a proof of the second part of the theorem.

22.6 Amitsur-Levitzki Theorem and the cohomology of $\mathfrak{gl}(n)$

We will sketch the proof of the Amitsur-Levitzki Theorem via this Lie algebra cohomology story. To simplify notation, we work over $\mathbb{C}$.

Notice that the identity S_{2n} holding in $\mathrm{Mat}_n(k)$ is equivalent to

$$\mathrm{Tr}(S_{2n+1}(x_1, \ldots, x_{2n+1})) = 0$$

for all $x_1, \ldots, x_{2n+1} \in \mathrm{Mat}_n(k)$. To see this, note that trace is cyclically invariant ($\mathrm{tr}(abc) = \mathrm{tr}(cab)$, etc.), so for each monomial in S_{2n+1}, we can cyclically permute the variables until x_1 is at the left. Factoring x_1 out, we obtain $\mathrm{Tr}(x_1 S_{2n}(x_2, \ldots, x_{2n+1})) = 0$. Since the trace pairing is nondegenerate, this implies that $S_{2n}(x_2, \ldots, x_{2n+1}) = 0$.

Now view $\mathrm{Mat}_n(\mathbb{C})$ as a Lie algebra, so $\mathfrak{gl}_n(\mathbb{C})$. The multilinear functional

$$(x_1, \ldots, x_{2i-1}) \mapsto \mathrm{Tr}(S_{2i-1}(x_1, \ldots, x_{2i-1}))$$

defines an element

$$\sigma_i = \sigma_{i,n} \in \Lambda^{2i-1}\mathfrak{g}^*$$

invariant under conjugation by $G = \mathrm{GL}_n(\mathbb{C})$, so $\sigma_i \in (\Lambda^{2i-1}\mathfrak{g}^*)^G$. For $G = \mathrm{GL}_n(\mathbb{C})$ and other complex reductive groups, there are isomorphisms

$$(\Lambda^\bullet \mathfrak{g}^*)^G \cong H^\bullet(\mathfrak{g}) \cong H^\bullet(G, \mathbb{C}) \cong H^\bullet(K, \mathbb{C}).$$

Here $H^\bullet(\mathfrak{g})$ is the **Lie algebra cohomology**, i.e. $\mathrm{Ext}^\bullet_{U(\mathfrak{g})}(\mathbb{C},\mathbb{C})$ (parallel to the definition of group cohomology). $H^\bullet(G,\mathbb{C})$ is the cohomology of G viewed as a topological space, while $K \subset G$ is a maximal compact subgroup and $H^\bullet(K,\mathbb{C})$ is the cohomology for K viewed as a topological space. For $G = \mathrm{GL}_n(\mathbb{C})$, the maximal compact subgroup K is the group $U(n)$ of unitary matrices, and

$$H^*(U(n),k) = \Lambda[c_{1,n}, c_{2,n}, \dots, c_{n,n}],\ \deg(c_{i,n}) = 2i-1.$$

This is graded and skew-commutative so $c_i^2 = 0$. This follows from induction and the fact that $U(n)/U(n-1) = S^{2n-1}$ (the $(2n-1)$-dimensional sphere). The restriction map

$$H^\bullet(\mathfrak{gl}(n)) \to H^\bullet(\mathfrak{gl}(n-1))$$

sends $c_{i,n} \mapsto c_{i,n-1}$ when $i \leqslant n-1$ and $c_{n,n} \mapsto 0$.

This gives a proof of the Amitsur-Levitski Theorem as follows:

We want to show that $\sigma_{i,n} = 0$ for $i > n$. We induct on n, so assume $\sigma_{i,n-1} = 0$ for $i > n-1$. So in particular

$$\sigma_{n+1,n} \in \ker(H^{2n+1}(\mathfrak{gl}(n)) \to H^{2n+1}(\mathfrak{gl}(n-1))).$$

We claim this map is injective: the kernel of the restriction map $H^\bullet(\mathfrak{gl}(n)) \to H^\bullet(\mathfrak{gl}(n-1))$ is generated by an element of degree $2n-1$ and $H^2(\mathfrak{gl}(n)) = 0$, so there is nothing in the kernel in degree $2n+1$. So $\sigma_{n+1,n} = 0$.

It remains to show that $\sigma_{i,n} = 0$ for $i > n+1$. The vanishing of $\sigma_{i,n}$ is equivalent to S_{2i} being an identity in $\mathrm{Mat}_n(\mathbb{C})$. But if the identity S_m holds, so does S_p for $p > m$ because one can sum over the symmetric group Σ_p by first summing over the Σ_m-cosets in Σ_p. This completes the induction.

Remark 22.12: $H^\bullet(\mathfrak{gl}(n))$ is in fact freely generated by $\sigma_{1,n}, \dots, \sigma_{n,n},$.

23.1 Proof of Amitsur-Levitski Theorem

We now give a self-contained proof of the first part of Theorem 22.9 using Cayley-Hamilton, due to Rossett.

Theorem 23.1 (Cayley-Hamilton): Let $x \in \mathrm{Mat}_n(R)$ be an $n \times n$ matrix with coefficients in R and $P_x(t) \in R[t]$ be its characteristic polynomial. Then $P_x(x) = 0$.

Proof (Sketch). The "easiest to remember" proof is to first reduce to $R = \mathbb{Z}$ by noting that the entries $P_x(x) = 0$ will be polynomials in the entries of x with integer coefficients. Then it suffices to show this for $R = \mathbb{C}$. But over $\mathbb{C}$ all matrices can be put in Jordan normal form, and for such a matrix $P_x(x) = 0$. □

Proof. A more aesthetically appealing proof: every matrix A has an adjoint matrix B such that $AB = BA = \det(A) \cdot I_n$ where I_n is the identity $n \times n$ matrix. Then in $\mathrm{Mat}_n(R[t])$, letting $A = t \cdot I_n - x$, by definition $\det(A) = P_x(t)$ and there exists B such that

$$BA = AB = P_x(t) \cdot I_n.$$

Let $\mathrm{Mat}_n(R[t]) = \mathrm{Mat}_n(R)[t]$ act on $\mathrm{Mat}_n(R)$ where $\mathrm{Mat}_n(R)$ acts via left multiplication and t acts by right multiplication by x. Then

$$A.I_n = 0 \Longrightarrow (P_x(t) \cdot I_n).I_n = 0$$

so $P_x(x) = 0$. □

Corollary 23.2: If $P_x(t) = t^n$, then $x^n = 0$.

Proof (of Theorem 22.9). It suffices to consider $R = \mathbb{Z}$ since multilinear identities are inherited by the extension of scalars. Since $\mathrm{Mat}_n(\mathbb{Z}) \subset \mathrm{Mat}_n(\mathbb{Q})$, it is enough to consider $R = \mathbb{Q}$. We will show for a certain matrix x that $\mathrm{Tr}(x^i) = 0$ for $i = 1, \dots, n-1$, which will imply that $P_x(t) = t^n$.

Consider an auxiliary ring $\Lambda = \Lambda^\bullet(\mathbb{Q}^{2n})$ (exterior algebra of $\mathbb{Q}^{2n}$) where $\mathbb{Q}^{2n}$ has basis $\varepsilon_1, \dots, \varepsilon_{2n}$. For $x_1, \dots, x_{2n} \in \mathrm{Mat}_n(\mathbb{Q})$ let

$$\boldsymbol{x} = \varepsilon_1 x_1 + \cdots + \varepsilon_{2n} x_{2n} \in \mathrm{Mat}_n(\Lambda) = \mathrm{Mat}_n(\mathbb{Q}) \otimes_{\mathbb{Q}} \Lambda.$$

Amitsur-Levitski will hold iff

$$\boldsymbol{x}^{2n} = \varepsilon_1 \wedge \cdots \wedge \varepsilon_{2n} \cdot S_n(x_1, \dots, x_n) = 0.$$

Notice that the even-degree wedges Λ^{ev} form a commutative ring, so Cayley-Hamilton applies here! Decompose

$$\Lambda^{\mathrm{ev}} = \mathbb{Q} \oplus \Lambda^2 \oplus \cdots \oplus \Lambda^{2n}.$$

So it remains to check that the coefficients of the characteristic polynomial of $\boldsymbol{x}^2$ vanish, i.e. $\mathrm{Tr}(\boldsymbol{x}^{2i}) = 0$. But this is true because

$$\mathrm{Tr}(x_1 \cdots x_{2i}) = \mathrm{Tr}(x_{2i} x_1 \cdots x_{2i-1})$$

and this cycle is an odd permutation because the number of letters is even. □

23.2 Primitive algebras and Kaplansky's theorem

Theorem 23.3 (Kaplansky): Let A be a primitive PI algebra over a field k with a homogeneous identity of degree d. Then it is simple of degree $m \leqslant d/2$ over its center, which is a (possibly different) field K.

Proof. Let L be a faithful simple module over A and $D := \mathrm{End}_A(L)$.
First, $\dim_D(L) \leqslant d/2$. If not, the image of $A \subset \mathrm{End}(L)$ would contain $\mathrm{Mat}_n(D)$ as a subquotient for $n > d/2$ (pick some collection of $n > d/2$ linearly independent vectors in L and consider the subalgebra A' of A that preserves vector space generated by these vectors, it then follows from the density theorem 3.3 that A' surjects onto $\mathrm{Mat}_n(D)$), which contradicts the easy part of Amitsur-Levitski.
Then we claim that D is finite over its center K. If not, pick a maximal commutative subfield $F \subset D$, which exists by Zorn's lemma. As we discussed earlier, WLOG we can assume the identity is multilinear, so it's inherited by extension of scalars and also holds in $F \otimes_K D$. By Azumaya-Nakayama, $F \otimes_K D$ is a simple ring. Moreover, by having D act by left multiplication and F by right multiplication, we get an action of $F \otimes_K D$ on D, and D is a simple module over $F \otimes_K D$ with

$$\mathrm{End}_{F \otimes_K D}(D) = Z_D(F) = F.$$

So by the argument in the previous paragraph, D is finite-dimensional over F. $F \otimes_K D \subset \mathrm{End}_F(D)$ is finite over K because $F \otimes_K D$ simple impiles D is a faithful $F \otimes_K D$-module. Thus, D is finite over K.
Finally, to get the degree bound, let E be a splitting field of D. Then

$$\mathrm{Mat}_n(D) \otimes_K E = \mathrm{Mat}_{n \cdot \deg D}(E) \Longrightarrow 2n \deg(D) \leqslant d$$

via the easy part of Amitsur-Levitski. □

23.3 Prime PI algebras and Posner theorem

Theorem 23.4 (Posner): Let A be a prime PI algebra. Then its center $Z = Z(A)$ is a domain. Moreover, $A \otimes_Z \mathrm{Frac}(Z) \cong \mathrm{Mat}_n(D)$ for some skew field D that is finite-dimensional over $K = \mathrm{Frac}(Z)$.

The proof follows from another fact about semi-prime PI algebras, which follows from Kaplansky's Theorem.

Theorem 23.5 (Rowen): Let A be a semi-prime PI algebra. Then every nonzero two-sided ideal meets the center.

Corollary 23.6: A prime PI ring A whose center is a field K is a central simple algebra over K.

Proof. By Rowen's theorem, every nonzero two-sided ideal in A meets Z. Thus, A is simple, and Kaplansky's theorem shows A is finite-dimensional over K. □

Proof (of Theorem 23.4). Z is a domain, since if $z_1 z_2 = 0$ for central elements $z_1, z_2 \in Z$, then $Az_1 \cdot Az_2 = 0$, contradiction. Homogeneous polynomial identities are inherited by extension of scalars, so $A \otimes K$ is simple by the Corollary 23.6. □

23.4 Central polynomials

The proof of Rowen's theorem is via central polynomials, which are noncommutative polynomials that take values in the center. We will be interested in Razmyslov's central polynomials, which, when you plug in $n \times n$ matrices, give back a scalar matrix (and are not identically zero).

First, we start with a linear algebra construction. Recall that $M = \mathrm{Mat}_n(k)$ has a nondegenerate trace pairing

$$\langle x, y\rangle = \mathrm{Tr}(xy).$$

This corresponds to the element $\tau \in M \otimes M$, $\tau = \sum_i m_i \otimes m_i^*$ (coevaluation) where the m_i and m_i^* are dual bases.

Lemma 23.7: For $x \in \mathrm{Mat}_n(k)$, $\sum_i m_i x m_i^* = \mathrm{Tr}(x) I_n$.

Proof. First let us check the identity for the standard basis $e = \{e_{ab}\}$ and $e^* = \{e_{ba}\}$, as $\mathrm{Tr}(e_{ab}e_{cd}) = \delta_{a=d} \cdot \delta_{b=c}$. Now we write $x = \sum_{c,d} x_{cd} e_{cd}$ and compute that

$$\sum_{a,b} e_{ab} \left(\sum_{c,d} x_{cd} e_{cd}\right) e_{ba} = \sum_{a,b} e_{ab} x_{bb} e_{bb} e_{ba} = \left(\sum_b x_{bb}\right)\left(\sum_a e_{aa}\right) = \mathrm{Tr}(x) \cdot I_n.$$

Now for an arbitrary m and m^*, then $m = Pe$ and $m^* = e^* Q$ for some $n^2 \times n^2$ matrices P, Q. But we get that

$$I_{n^2} = \langle m, m^*\rangle = P\langle e, e^*\rangle Q = PQ,$$

so $PQ = I_{n^2}$ are inverse matrices. Therefore

$$\sum_i m_i x m_i^* = \sum_{i,j,k} p_{ij} e_j x e_k^* q_{ki} = \sum_{i,j,k} q_{ki} p_{ij} e_j x e_k^* = \sum_{i,j,k} \delta_{k=j} e_j x e_k^* = \sum_j e_j x e_j^* = \mathrm{Tr}(x) \cdot I_n.$$

□

Now we will look for noncommutative polynomials in n^2 variables such that $P(m_1, \ldots, m_{n^2})$ is an element of the dual basis.

Definition 23.8: The **Capelli polynomial** is

$$C(x_1, \ldots, x_N, y_0, \ldots, y_N) = \sum_{\sigma \in \Sigma_n} (-1)^{|\sigma|} y_0 x_{\sigma(1)} y_1 \cdots x_{\sigma(N)} y_N.$$

It's like S_N, but it inserts "separating" variables y_i.

Lemma 23.9: Let a_i be a basis of $\mathrm{Mat}_n(k)$ and a_i^* be its dual basis (w.r.t. the trace pairing). Define $C_i := \tau_i(C)$ where τ_i is a linear endomorphism on the space of multilinear noncommutative polynomials defined by $\tau_i(u x_i v) = vu$. Then for $b_0, \ldots, b_{n^2}$ any matrices,

$$\mathrm{Tr}(C(a_k, b_l)) a_i^* = C_i(a_k, b_l).$$

Proof. Let $t = \mathrm{Tr}(C(a_k, b_l))$. We need to prove that

$$\mathrm{Tr}(a_j C_i(a_k, b_l)) = \begin{cases} 0, & i \neq j \\ t, & i = j \end{cases}.$$

Using that $\mathrm{Tr}(a_j(vu)) = \mathrm{Tr}(u a_j v)$, each monomial recovers the trace from the corresponding τ_i (i.e., $\mathrm{Tr}(a_j \tau_i(C)) =$

$\operatorname{Tr} C(..., x_i = a_j, ...))$. When $i \neq j$, we get

$$\operatorname{Tr}(C(a_1, \dots, a_{i-1}, a_j, a_{i+1}, \dots, b_l)) = 0$$

(here we have replaced a_i with a_j and used that C is antisymmetric to conclude that the resulting trace is zero). When $i = j$, we get t. Moreover, plugging in the elementary matrices, we get that t is not uniformly zero. □

Definition 23.10: The **Razmyslov polynomial** is

$$Z_n(x_1, \dots, x_{n^2}, y_0, \dots, y_{n^2}, z) = \sum_i x_i z C_i(x_1, \dots, x_{n^2}, y_0, \dots, y_{n^2}).$$

Theorem 23.11: The map $\operatorname{Mat}_n(k)^{2n^2+2} \longrightarrow \operatorname{Mat}_n(k)$ sending $x_1, \dots, x_{n^2}, y_0, \dots, y_{n^2}, z$ to $Z_n(x_k, y_l, z)$ takes values in the scalar matrices and is not identically zero.

Proof. The previous lemmas tell us that

$$Z_n(x_k, y_l, z) = \operatorname{Tr}(C(x_k, y_l)) \operatorname{Tr}(z) I_n.$$

So we should find x_k, y_l for which $\operatorname{Tr}(C(x_k, y_l)) \neq 0$. Let $m_1, \dots, m_{n^2}$ be the n^2 elementary matrices. Now if we have $m_a = e_{kl}$ and $m_{a+1} = e_{k'l'}$, let

$$y_a := e_{lk'},\ a = 1, \dots, n^2 - 1$$

and saying $m_1 = e_{kl}$ and $m_{n^2} = e_{l'1}$, let

$$y_0 := e_{l'k}.$$

Now setting $x_k = m_k$, the monomial corresponding to the identity permutation evaluates to e_{11} while all the other monomials evaluate to 0.
Therefore, Z_n is nonzero here. □

23.5 Rowen's Theorem for semi-primitive algebras

To prove Rowen's Theorem, first we prove a version for semi-primitive PI algebras that uses the central polynomials above.

Proposition 23.12: Let A be a semi-primitive PI algebra. Then every nonzero two-sided ideal meets the center.

Proof. Let L be an irreducible A-module. Then for $D = \operatorname{End}_A(L)$ and $K = Z(\operatorname{End}_A(L))$, Kaplanksy's theorem, gives us a bound on $d(L) := \dim_D(L) \cdot \deg(D)$. The image $\bar{A}_L$ of the map $A \longrightarrow \operatorname{End}(L)$ is isomorphic to $\operatorname{Mat}_m(D)$, and choosing a splitting field F of D, we get

$$\bar{A}_L \otimes_K F \cong \operatorname{Mat}_{d(L)}(F).$$

Let n be the maximal $d(L)$ such that $I \not\subset \operatorname{Ann}(L)$. Then we claim that our central polynomial $c = Z_n(x_k, y_l, z)$ for $z \in I$ lies in the center of A. We show that it will go to a central element in any irreducible L; this is enough because A is semi-primitive. If $d(L) > n$, then c acts by zero in L because z does. If $d(L) < n$, then c also maps to zero because Z_n is an identity in $\operatorname{Mat}_m(k)$ for $m < n$. If $d(L) = n$, z becomes a scalar matrix after extending scalars to F as above, so it lands in K.
The last thing is to show that $c \neq 0$. To do so, pick L with $d(L) = n, IL \neq 0$. Then I maps onto $\bar{A}_L = \operatorname{Mat}_m(D)$ so it suffices to show that Z_n is not an identity in $\operatorname{Mat}_m(D)$. But since identities are preserved by extension of scalars and Z_n is not an identity in $\operatorname{Mat}_n(F)$, Z_n is nonzero in $\operatorname{Mat}_m(D)$. □

23.6 Proof of Rowen's Theorem (for real)

Rowen's theorem follows from the above weaker version and

Theorem 23.13: If R is a semi-prime PI algebra, then $R[t]$ is a semi-primitive PI algebra.

Proof (of Theorem 23.5). If R is a semi-prime PI algebra then $R[t]$ is a PI algebra since extension of scalars to $R[t]$ preserves multi-linear identities, hence the PI property. Now if $I \subset R$ is a nonzero ideal, then $I[t] \subset R[t]$ will meet the center of $R[t]$ (by Theorem 23.13 and Proposition 23.12). But if $I[t]$ meets the center, then so does I. □

Theorem 23.13 will follow from the following.

Definition 23.14: A **nil ideal** is an ideal consisting of nilpotent elements.

Theorem 23.15 (Amitsur): If R has no nonzero nil ideals, then $R[t]$ is semi-primitive.

Proposition 23.16: A semi-prime PI algebra contains no nil ideals.

Proof (of Theorem 23.15). Let $J \subset R[t]$ be the Jacobson radical and suppose $p(t) = \sum a_i t^{r_i} \in J$ is a nonzero element of the Jacobson radical. WLOG, we can assume that the length of this sum is the minimal possible for a nonzero $p \in J$. Then the a_i must pairwise commute; otherwise, $[a_i, p]$ will be a shorter nonzero polynomial in J.
This implies that $1 + tp(t)$ is invertible in $R[t]$, and the coefficients of $(1 + tp(t))^{-1} \in R[[t]]$ lie in the commutative subring $S \subset R$ generated by the a_i. But for a commutative ring S, $1 + tp(t) \in S[t]$ is invertible iff all its coefficients are nilpotent: otherwise, we could find a maximal ideal $\mathfrak{m} \subset S$ such that $p \notin \mathfrak{m}[t]$ and p would be invertible over $(S/\mathfrak{m})[t]$, but nonconstant polynomials over fields cannot be invertible.
Therefore, each a_i is nilpotent. Then the set of all a_1 such that $q(t) = \sum a_i t^{r_i} \in J$ for some $a_2, \dots, a_n$ is a nil ideal in R. So $J = 0$. □

Finally, we prove Proposition 23.16.

Lemma 23.17: Suppose that a ring R satisfies the ascending chain termination condition for right annihilators. If R is semi-prime, then every nil left ideal is zero.

Proof. Suppose I is a nil left ideal. WLOG we can assume that $I = Ra$ for some $a \in R$. Let $J = \mathrm{rAnn}(b) \subsetneq R$ for $b \in I, b = xa$ be maximal among right annihilators of (nonzero) elements in I. Then if $b^n \neq 0$, then

$$\mathrm{rAnn}(b^n) \supset \mathrm{rAnn}(b) \Rightarrow \mathrm{rAnn}(b^n) = \mathrm{rAnn}(b).$$

Hence $b^2 = 0$; otherwise for $n \geqslant 2$ we have $b^n \neq 0$ and $b^{n+1} = 0$, so $b \in \mathrm{rAnn}(b^n)$ and $b \in \mathrm{rAnn}(b)$.
We also claim $bRb = 0$. To see this, fix $y \in R$ and consider $c = yb \neq 0$. Pick n such that $c^n \neq 0, c^{n+1} = 0$. Then

$$c \in \mathrm{rAnn}(c^n),\ \mathrm{rAnn}(c^n) \supset \mathrm{rAnn}(b) \Rightarrow \mathrm{rAnn}(c^n) = \mathrm{rAnn}(b).$$

Then $c \in \mathrm{rAnn}(b) \Rightarrow byb = 0$. Thus, RbR is a nonzero nilpotent ideal, contradicting that R is semi-prime. □

Lemma 23.18: A prime PI ring satisfies the ascending chain termination condition for right and left annihilators.

Proof. Suppose $P(x_1, \dots, x_n) = \sum a_\sigma x_{\sigma(1)} \cdots x_{\sigma(n)}$ is a multilinear identity holding in R and $I_1 \subsetneq I_2 \subsetneq \cdots$ is an infinite ascending chain of left annihilators. Let $J_i = \mathrm{rAnn}(I_i)$, so $I_i \subset \mathrm{lAnn}(J_i)$.
Now evaluate P at $x_i \in I_i$. WLOG we can assume P is the smallest degree of an identity holding for such a choice of variable values. So for $y \in J_{n-1}$, $P(x_1, \dots, x_n)y = 0$. Also $x_{\sigma(1)} \cdots x_{\sigma(n)} y = 0$ when $\sigma(n) \neq n$ since $y \in J_{n-1}$, $x_{\sigma(n)} \in I_{\sigma(n)}$, and $I_{\sigma(n)} J_{n-1} = 0$.
So we can write $P = Q(x_1, \dots, x_{n-1})x_n +$ monomials not ending in x_n. The previous paragraph shows that for any $x_i \in I_i$,

$$Q(x_1, \dots, x_{n-1}) I_n J_{n-1} = 0.$$

But R is prime, so $I_n J_{n-1} \neq 0$. Hence $Q(x_1, \dots, x_{n-1}) = 0$, which contradicts the degree minimality assumption. □

We've now proved Rowen's theorem if R is prime. If R is semi-prime, we need this last lemma:

Lemma 23.19: A semi-prime ideal is an intersection of prime ideals.

Proof. Let $I \subset R$ be a semi-prime ideal, i.e., $\bar{R} := R/I$ is semi-prime. Let $a \in \bar{R}$, $a \neq 0$. Since $\bar{R}$ is semi-prime, there exists $a_1 = a$ and $a_{i+1} = a_i x_i a_i$ such that $a_i \neq 0$ (construct a_i inductively, using that $(\bar{R}a_i\bar{R})^2 \neq 0$). Let J be a maximal ideal in $\bar{R}$ not containing a_i for any $i \geqslant 1$; this exists by Zorn's lemma.
Suppose J is not prime (i.e. ring $\bar{R}/J$ is not prime), so $x\bar{R}y \subset J$ for some $x, y \notin J$. Then for some i, $a_i \in \bar{R}x\bar{R} + J$ and $a_i \in \bar{R}y\bar{R} + J$ (use that if a_n lies in some ideal then a_k for every $k \geqslant n$ lies in the same ideal). But then

$$a_{i+1} \in \bar{R}x\bar{R}y\bar{R} + J = J$$

which contradicts our choice of J. So $a \notin J$ and J is prime.
Since we can find such J for each $a \neq 0$, then I is an intersection of these prime ideals. □

Therefore, a semi-prime ring R can be realized as a subring in the product of prime rings R/J. If R is a PI ring, then each R/J is such, so a nil ideal in R has zero image in R/J for all J, thus, it is zero. This finally proves Rowen's theorem.

24 May 11 - Gelfand-Kirillov dimension

24.1 Growth of algebras and Gelfand-Kirillov dimension

Recall that we defined the growth of a (finitely generated) algebra as follows: pick a finite-dimensional space of generators V, which gives us a (surjective) homomorphism $TV \twoheadrightarrow A$ and induces a filtration of A by setting

$$A_{\leqslant n} = \operatorname{im}\left(\bigoplus_{i \leqslant n} V^{\otimes i}\right).$$

Let

$$d(n) := \dim(A_{\leqslant n}).$$

Then we saw that the order of growth was independent of V. Recall that A has

- subexponential growth if $d(n) < cn^\alpha$ for some c for all $\alpha > 1$
- exponential growth if $\limsup \sqrt[n]{d(n)} > 1$
- polynomial growth if there exists c, δ such that $d(n) \leqslant cn^\delta$.

Definition 24.1: The **Gelfand-Kirillov dimension** of an algebra A is

$$\inf\{\delta \mid \exists c, d(n) \leqslant cn^\delta\}.$$

That is, $\operatorname{GKdim}(A) \in \mathbb{R}_{\geqslant 0} \cup \{\infty\}$, and this is well-defined since another function d' such that there exists $a \geqslant 1$ with

$$d'(n/a) \leqslant d(n) \leqslant d'(an)$$

will lead to the same value.

Remark 24.2: There are similar definitions for finitely generated groups; if they have exponential growth, then they're called hyperbolic groups. (There is not much similarity in the methods and theorems though.)

Remark 24.3: If $A_{\leqslant n+1} = A_{\leqslant n}$, then $A = A_{\leqslant n}$ and $d(n)$ will eventually be $\dim A$. Hence, if $\dim A < \infty$, $\operatorname{GKdim}(A) = 0$. Otherwise, we'll always have $d(n) \geqslant n + 1$, so the GK dimension will be $\geqslant 1$.

Lemma 24.4:
a) If $\operatorname{GKdim}(A) < 1$, then A is finite-dimensional so $\operatorname{GKdim}(A) = 0$.
b) $\operatorname{GKdim}(A[t]) = \operatorname{GKdim}(A) + 1$
c) $\operatorname{GKdim}(A[a^{-1}]) = \operatorname{GKdim}(A)$ if a is central and regular.

Proof. a) If $d(n+1) = d(n)$ for some n, then $A_{\leqslant n+1} = A_{\leqslant n}$. So $A = A_{\leqslant n}$ and $d(n)$ will eventually be $\dim A$. Hence, if $\dim A < \infty$, $\mathrm{GKdim}(A) = 0$. Otherwise, we'll always have $d(n+1) \geqslant d(n) + 1$, so $d(n) \geqslant n$ and the GK dimension will be $\geqslant 1$.

b) Let $B = A[t]$ and $V_B = V_A \oplus kt$. Then $B_{\leqslant n} = t^n A_{\leqslant 0} \oplus t^{n-1} A_{\leqslant 1} \oplus \cdots \oplus A_{\leqslant n}$, so $\dim B_{\leqslant n} \leqslant (n+1) \dim A_{\leqslant n}$. Also, $\dim B_{\leqslant n} \geqslant n \dim A_{\leqslant n}$, so $\mathrm{GKdim}(A[t]) = \mathrm{GKdim}(A) + 1$.

c) Again, add a^{-1} to the space of generators $V_A \ni a$, so that $V_B = V_A + k \cdot a^{-1}$. Then $\dim(A_{\leqslant n}) \leqslant \dim(B_{\leqslant n}) \leqslant \dim(A_{\leqslant 2n})$; the first inequality is because $A_{\leqslant n} \subset B_{\leqslant n}$, and the second inequality is because $B_{\leqslant n} \hookrightarrow A_{\leqslant 2n}$ via multiplication by a^n: the elements of $B_{\leqslant n}$ look like $a^{-k}\alpha$ for $\alpha \in A_{\leqslant n-k}$ (since a is central), so $a^n \cdot a^{-k}\alpha = a^{n+k}\alpha \in A_{\leqslant 2n}$. So $\mathrm{GKdim}(A[a^{-1}]) = \mathrm{GKdim}(A)$.

□

Example 24.5: Part b) implies that the GK dimension of $k[x_1, \ldots, x_n]$ is n.

24.2 Warfield's Theorem

The GK dimension of a noncommutative ring can take any value $\geqslant 2$.

Theorem 24.6 (Warfield): For any real $\delta \geqslant 2$, there exists an algebra with 2 generators whose GK dimension is δ.

Proof. Part b) of Lemma 24.4 implies we only have to show this for $\delta \in (2,3)$. We will construct a quotient of $k\langle x, y\rangle$ by monomials. Fix a monotonically increasing sequence γ_n, $n = 1, \ldots,$, let $I \subset k\langle x, y\rangle$ be the ideal spanned by the monomials of degree at least 3 in y and the monomials

$$x^i y x^j y x^k,\ j < \gamma_n, n = i + j + k.$$

The quotient $k\langle x, y\rangle / I$ is a graded algebra A; let A_n be the component of degree n. Then

$$\dim(A_n) = 1 + n + \binom{n + 2 - \gamma_n}{2}$$

where the $1 + n$ comes from monomials of degree $0, 1$ in y.
If we take $q \in (0,1)$ and set $\gamma_n = n - [n^q]$, then $\mathrm{GKdim}(A) = \max(2, 2q+1)$. Hence this gives you anything in $(2,3)$. □

Remark 24.7: This doesn't happen for finitely presented monomial algebras. Notice that for every finitely generated algebra A, one can find a finitely generated monomial algebra $\bar{A}$ with the same growth function by setting $\bar{A}$ to be the associated graded for a filtration on A. But the same construction for A finitely presented does not imply that $\bar{A}$ is finitely presented.

24.3 GK dimension of path algebras

Let Γ be a finite oriented graph, and let its edges, the arrows, be labeled $x_1, \ldots, x_m$. The *path algebra* A is the graded algebra with basis the oriented paths in Γ. The product uv of two paths is defined to be the composed path $u \circ v$ if the paths can be composed, i.e., if v starts where u ends, and 0 if the paths can not be composed.

The algebra is graded by the *length* of a path, the number of its arrows: A_n is the vector space spanned by paths of length n.

It is customary to adjoin paths of length zero which represent the vertices of the graph. They are idempotent elements in the path algebra A, and their sum is 1.

For example, the paths in the graph $\circ \xrightarrow{x_1} \circ \xrightarrow{x_2} \cdots \xrightarrow{x_m} \circ$ are the words $x_i x_{i+1} \cdots x_j$ for $1 \leq i \leq j \leq m$. This path algebra is the algebra of upper triangular $(m+1) \times (m+1)$ matrices, with $x_i = e_{i-1,i}$.

The path algebra is finite dimensional iff the (finite) graph Γ has no oriented loops, or cycles. But a cycle gives us a path, say u, whose powers u^n are all distinct and nonzero. The next proposition shows how the GK dimension of the path algebra can be read off from the geometry of the graph.

Proposition 24.8: Let A be the path algebra of a finite oriented graph Γ.

(i) If Γ contains two cycles which have a vertex in common, then A has exponential growth.

(ii) Suppose that Γ contains r cycles, none of which have a vertex in common. Then $GK\dim(A) \leq r$. Moreover, $GK\dim(A) = r$ if and only if there is a path that traverses all of the cycles.

Proof sketch. (i) Say that a vertex is in common to two distinct cycles, and let u, v be the paths which traverse these loops, starting and ending at a common vertex. Then the words in u, v represent distinct paths. In fact, they represent distinct elements in the fundamental group of the graph. So A contains the free ring $k\langle u, v\rangle$.

(ii) Suppose for instance that there are exactly two distinct cycles, and say that u, v are paths which traverse them. There may or may not be some paths y connecting u to v, i.e., such that $uyv \neq 0$. But if such a path exists, then there can be no path in the opposite direction, because if $vzu \neq 0$, then yz would be a cycle having a vertex in common with u. If y exists, then the paths which can be built using u, y, v are $u^i y v^j$. Since we cannot return to u from v, every path has the form $wu^i y v^j w'$, where each of the subpaths w, w', y is a member of a finite set. This leads to very regular quadratic growth.

□

24.4 Bergman gap theorem

The proof of this theorem is presented in [12] chapter 2.

Theorem 24.9 (Bergman gap): There is no finitely generated algebra whose GK dimension is strictly between 1 and 2.

Proof. The theorem follows from the below proposition. To reduce to a graded algebra generated in degree 1, we can reduce to $A = k\langle x_1, \dots, x_n\rangle/J$ where J is a monomial ideal. Then take the associated graded, first by total degree, then by lexicographical order. Then either $\dim A_d \geqslant d$, which implies that the GK dimension is at least 2, or there exists d with $\dim A_d < d$, which (by the proposition) implies that $\dim A_n$ is bounded by a constant and the GK dimension is at most 1. □

Proposition 24.10: If A is a graded algebra generated in degree 1 and there exists d such that $\dim A_d < d$, then $\mathrm{GKdim}(A) \leqslant 1$.

Proof. WLOG we can assume that all relations are monomial in degree d. To prove this, we define "allowed words", where a word is allowed iff all its subwords of length d are allowed. Let S be the set of allowed words of degree d and suppose $|S| \leqslant d$. Then the number of allowed words of degree N is bounded.
This reduces to

Lemma 24.11: Assume that there at most d allowed words of length d. Then for $n \geqslant 2d$, every allowed word of length n has the form $w = w_1 w_2 w_3$ where w_2 is p-periodic for $p \leqslant d$, $|w_1|, |w_3| \leqslant d - p$, and $|w_2| \geqslant d + p$. (A finite word $x_1 \cdots x_n$ is p-periodic if $x_{i+p} = x_i$ when $i, i + p \in [1, \dots, n]$.)

Proof. We induct on $|w|$. The base case is $|w| = 2d$. Such a word will have $d + 1$ subwords of length d, but since there are only d distinct allowed words, at least two of these coincide and we have the desired periodicity. Now we need the following:

Lemma 24.12: If a periodic word with minimal period p contains two equal subwords of length $\geqslant p - 1$, then they are np letters apart.

Proof. Extend the word to an infinite p-periodic word. Suppose the equal subwords are $x_{i+1} \cdots x_{i+r}$ and $x_{j+1} \cdots x_{j+r}$ with $r \geqslant p - 1$. Then the subwords $x_i x_{i+1} \cdots x_{i+p-1}$ and $x_j x_{j+1} \cdots x_{j+p-1}$ are each a full period of the word x. Since $x_{i+q} = x_{j+q}$ for all $1 \leqslant q \leqslant r$, then $x_i = x_j$ also.
So x also has equal subwords $x_i \cdots x_{i+p-1}, x_j \cdots x_{j+p-1}$. Let the word have length m and $1 \leqslant \ell \leqslant m$, and let t be an integer such that $\ell + tp = i + s$ for $0 \leqslant s \leqslant p - 1$. Then

$$x_{\ell+(j-i)} = x_{\ell+(j-i)+tp} = x_{i+s+(j-i)} = x_{j+s} = x_{i+s} = x_{\ell+tp} = x_\ell$$

so x has period $j - i$. Thus x has period equal to the greatest common divisor of $p, j - i$ and the minimality of p implies that $p | j - i$ as desired. □

Now we finish the proof of the lemma. Write $w = x_1 w'$ and $w' = w_1' w_2' w_3'$. If $|w_1'| < d - p$, there's nothing to do. Otherwise, in $x_1 x_2 \cdots x_d$, find two coinciding length d words. These intersect w_2 by at least $p - 1$, so their intersections with w_2' differ by a shift by n and $p|n$. One of them ends at x_{2d-p} (or to the left) so it contains x_{d-p+1}. Hence $x_{d-p+1} = x_{d-p+1+n} = x_{d+1}$. □

This finishes the proof of the proposition. □

24.5 Ufnarovskii graph

Another way of working with allowed words is via the overlap graph, called the Ufnarovskii graph; the proof of the theorem can also be interpreted via the graph. Consider an oriented graph U whose vertices are allowed length d words and which has an edge between w_1 and w_2 iff w_2 is obtained from w_1 by removing the first letter and adding a letter at the end. Then paths of length $n - d$ correspond to allowed words of degree $n \geqslant d$.

The proof of the Bergman gap theorem can be restated as follows: if there are at most d allowed words of length d, show that U contains at most one oriented cycle. Then any path in the graph enters the cycle at most once, traverses the cycle, then leaves the cycle; this is the factorization $w = w_1 w_2 w_3$ in the lemma above (see [1, Section VI.4]).

24.6 Smoktunowicz and Berele theorems

We state without proof two related results:

Theorem 24.13 (A. Smoktunowicz): The Gelfand-Kirillov dimension of a graded domain cannot fall within the open interval $(2, 3)$.

Theorem 24.14 (Berele): Finitely generated PI algebras have finite GK dimension.

24.7 GK dimension of a module

Definition 24.15: We can likewise define the Gelfand-Kirillov dimension of a finitely generated module over A by defining

$$d_M(n) = \dim M_{\leqslant n}$$

where we pick generators $W \subset M$ and $M_{\leqslant n} = A_{\leqslant n} \cdot W$, and setting

$$\mathrm{GKdim}(M) = \inf\{\delta \mid \exists c, d_M(n) \leqslant cn^\delta\}.$$

Again, this is not dependent on the choices of W.

Definition 24.16: We say that the GK dimension is **exact** for modules over an algebra A if for $M \supset N$,

$$\operatorname{GKdim}(M) = \max(\operatorname{GKdim}(N), \operatorname{GKdim}(M/N)).$$

Example 24.17: GK dimension is exact for finitely generated modules over Noetherian PI algebras.

Suppose that A is an algebra with commutative associated graded (which also is then automatically finitely generated, hence Noetherian). Then the GK dimension is exact for (f.g. modules over) A, because

Proposition 24.18: In this case, $\operatorname{GKdim}(M)$ is the dimension of the support of the $\operatorname{gr}(A)$ module $\operatorname{gr}(M) = \bigoplus M_{\leqslant d}/M_{\leqslant d-1}$.

In fact, there is a closer relation between the commutative and noncommutative pictures. Let $\operatorname{gr} A = \bar{A}$. Given an increasing filtration on A such that $\bar{A}$ is commutative, let a good filtration on M be a filtration such that $M = \bigcup M_{\leqslant d}$, $\bigcap M_{\leqslant d} = 0$, $A_{\leqslant 1} M_{\leqslant n} \subset M_{\leqslant n+1}$, and $\operatorname{gr} M = \bar{M}$ is a finitely generated $\bar{A}$ module.

Lemma 24.19: For A, M as above, the (set theoretic) support $\operatorname{supp}(\operatorname{gr} M) \subset \operatorname{Spec}(\bar{A})$ and does not depend on the choice of filtration. Moreover, the class of $\bar{M}$ in $K\left(\bar{A}\text{-mod}_S\right)$ (the Grothendieck group) is independent of the choice of the filtration, where $\bar{A}\text{-mod}_S$ is the category of finitely generated $\bar{A}$-modules with set-theoretic support contained in S.

Remark 24.20: The expression "set-theoretic support" refers to thinking of finitely generated $\bar{A}$-modules as coherent sheaves on $\operatorname{Spec}(\bar{A})$. Closed subsets $S \subset \operatorname{Spec}(\bar{A})$ correspond to radical ideals $I_S \subset \bar{A}$, and M is set-theoretically supported on S iff every element of M is annihilated by some power of I_S. Note that being scheme-theoretically supported on S would instead mean that M is annihilated by I_S, which is stronger.

Proof (of lemma, sketch). Given two good filtrations $M_{\leqslant d}$ and $M'_{\leqslant d}$, find m such that

$$M_{\leqslant d-m} \subset M'_{\leqslant d} \subset M_{\leqslant d-m+1}.$$

Inducting on M, we can reduce to the situation when $m = 0$ and

$$M_{\leqslant d} \subset M'_{\leqslant d} \subset M_{\leqslant d+1}.$$

Let

$$N = \bigoplus M_{\leqslant d}/M'_{\leqslant d-1}, \quad N' = \bigoplus M'_{\leqslant d}/M_{\leqslant d}.$$

Then there are short exact sequences

$$0 \to N' \to \bar{M} \to N \to 0$$
$$0 \to N \to \bar{M}' \to N' \to 0$$

which shows both statements. □

Remark 24.21: $\bar{M}$ is naturally graded, but the class of $\bar{M}$ in the Grothendieck group of graded $\bar{A}$-modules may depend on the choice of the filtration. This is because one can equip N, N' with a grading so that the first displayed SES is one of the graded modules, but the arrows in the second one will not agree with the grading.

24.8 Projective covers of graded modules

Let $A = k \oplus A_1 \oplus \cdots$ be a noetherian, connected graded algebra. The term connected just means that $A_0 = k$. In the next two sections we work primarily with graded right A-modules. By finite module we mean a finitely generated module. A *map* $\phi : M \longrightarrow N$ of graded modules is a homomorphism which sends $M_n \longrightarrow N_n$ for every n. The modules we consider will all be *left bounded*, which means that $M_n = 0$ if $n \ll 0$.

The *shift* $M(r)$ of a module M is defined to be the graded module whose term of degree n is $M(r)_n = M_{n+r}$. In other words $M(r)$ it is equal to M except that the degrees have been shifted. The reason for introducing these shifts is to keep track of degrees in module homomorphisms. For example, if $x \in A_d$ is a homogeneous element of degree d, then right multiplication by x defines a map of graded modules $A(r) \xrightarrow{\rho_x} A(r+d)$. Since all linear maps $A_A \longrightarrow A_A$ are given by left multiplication by A, this identifies the set of maps:

Corollary 24.22: $\mathrm{Hom}_{\mathrm{gr}}(A(r), A(s)) = A_{s-r}$.

If M is a graded right module and L is a graded left module, the tensor product $M \otimes_A N$ is a graded vector space, the degree d part of which is generated by the images of $\{M_n \otimes_k L_{d-n}\}$.

The symbol k will also denote the left or right A-module $A/A_{>0}$. It is a graded module, concentrated in degree zero, i.e., $k_0 = k$ and $k_n = 0$ for $n \neq 0$. For any module M, $MA_{>0}$ is a submodule, and

$$M \otimes k = M \otimes (A/A_{>0}) \cong M/MA_{>0}.$$

This is a graded vector space, and it is finite dimensional if M is finitely generated.

Proposition 24.23 (Nakayama Lemma):
(i) Let M be a left bounded module. If $M \otimes k = 0$, then $M = 0$.
(ii) A map $\phi : M \longrightarrow N$ of left bounded graded modules is surjective if and only if the map $M \otimes_A k \longrightarrow N \otimes_A k$ is surjective.

Proof. (i) Assume that M is not the zero module, and let d be the smallest degree such that $M_d \neq 0$. Then $(MA_{>0})_d = 0$, so

$$0 = (M \otimes k)_d \cong M_d/(MA_{>0})_d = M_d.$$

(ii) The second assertion follows from the right exactness of tensor product.

□

Definition 24.24: A map $P \longrightarrow M$ of finite graded modules is a **projective cover** of M if P is projective and if the induced map $P \otimes k \longrightarrow M \otimes k$ is bijective.

Proposition 24.25:
(i) Let $\phi : M \longrightarrow N$ be a surjective map of finite graded modules. If N is projective, then ϕ is bijective.
(ii) Every finite graded projective A-module is isomorphic to a finite direct sum of shifts of $A_A : P \cong A(r_i)$.
(iii) If $P' \longrightarrow P \longrightarrow M \longrightarrow 0$ is an exact sequence of finite graded modules with P', P projective, then P is a projective cover of M if and only if the map $P' \otimes k \longrightarrow P \otimes k$ is the zero map.

Proposition 24.26: Let

$$\mathcal{P} \longrightarrow M := \{\cdots \longrightarrow P_1 \longrightarrow P_0 \longrightarrow M \longrightarrow 0\}$$

be a projective resolution of a finite module M, and define M_i by $M_0 = M$ and $M_i = \ker(P_{i-1} \longrightarrow M_{i-1})$ for $i > 0$. The following conditions are equivalent. If they hold, the resolution is said to be a **minimal (projective) resolution**.
(a) P_i is a projective cover of M_i for all i,
(b) if $P_0 \longrightarrow M$ is a projective cover of M and for all $i > 0$, the induced maps $P_i \otimes k \longrightarrow P_{i-1} \otimes k$ are zero.

Corollary 24.27: Let $\mathcal{P} \longrightarrow M$ be a minimal projective resolution of a module M. Then $P_i \otimes k \cong \mathrm{Tor}_i^A(M, k)$.

Proof. The Tor are computed as the homology of the complex $\mathcal{P} \otimes k$. Since the maps in this complex are zero, $H_i(\mathcal{P} \otimes k) = P_i \otimes k$. □

Corollary 24.28: Let $\mathcal{P} \to k \to 0$ be a minimal projective resolution of the right module k, and say that

$$P_i \cong \bigoplus_j A(-r_{ij}).$$

The minimal projective resolution of k as left module has the same shape, i.e., the number of summands and the shifts r_{ij} which appear are the same.

Proof. $P_i \otimes k \cong \mathrm{Tor}_i^A(k,k)$, and $\mathrm{Tor}_i^A(k,k)$ can be computed using either a projective resolution of the left module k or a projective resolution of the right module k. □

Remark 24.29: Let $P = \bigoplus_i A(p_i)$ and $P' = \bigoplus_j A(q_j)$ be finite projective modules. Corollary 24.22 shows that

$$\mathrm{Hom}_{\mathrm{gr}}(P, P') = \bigoplus_{i,j} A(q_j - p_i).$$

The term $A(q_j - p_i)$ is zero unless $p_i \leq q_j$, because $A_n = 0$ if $n < 0$. If $\phi : P \to P'$ is a map, then $\phi \otimes k = 0$ is zero if and only if no entry ϕ_{ij} is a nonzero constant. This means $\phi_{ij} \in A_{q_j - p_i}$ is zero unless $p_i < q_j$.
Suppose that ϕ appears in a minimal projective resolution of some module. Then for every p_i, the summand $A(p_i)$ of P must have a nonzero image in P'. Together with the condition that $\phi \otimes k = 0$, this implies that p_i must be strictly less than at least one index q_j. So with the notation above, the indices $-r_{ij}$ are decreasing with i. However, because various shifts can appear, the overlapping indices confuse the situation.

24.9 Poincaré series

Let A be a noetherian connected graded algebra. The *Hilbert function* of A is the sequence $a_n = \dim_k A_n$. As we have seen, the Hilbert function is closely related to the growth of the algebra. We also consider the power series

$$h(t) = \sum_{n=0}^{\infty} a_n t^n,$$

which is called the *Hilbert series* of A.

Lemma 24.30:
(i) The radius of convergence r of the Hilbert series $h(t)$ is < 1 if and only if the Hilbert function has exponential growth.
(ii) For a sequence $a(n)$ of exponential growth, there exist $0 < r_1 < r_2 < \cdots$ such that

$$a(r_k) < \sum_{i=1}^{k-1} a(r_k - r_i).$$

Proof. (i) The root test tells us that $r = \limsup (a_n)^{1/n}$.
(ii) Having fixed $a(1), \ldots, a(m)$, there are infinitely many n such that

$$a(n) \geqslant \alpha^{r_i} a(n - r_i), i = 1, \ldots, m.$$

We can make the choice such that $\alpha^{r_k} > 2^k$.

□

Theorem 24.31 (Stephenson-Zhang): If A is right (or left) Noetherian, it has subexponential growth.

Proof. Apply the previous lemma to $a(n) = \dim A_n$. Inductively choose $x_i \in A_{r_i}$ such that $x_k \notin \sum_{i=1}^{k-1} x_i A_{k-i}$. □

Suppose that A has finite global dimension d. This means that every finite graded module has a graded projective resolution of length $\leq d$. Then one can obtain a recursive formula for the Hilbert function in terms of a resolution of the A-module k. (It is a fact that if k has a finite projective resolution, then A has finite global dimension, i.e, every finite A-module has a finite projective resolution, but never mind.)

Say that the minimal projective resolution is

$$0 \to P_d \xrightarrow{f_d} \cdots \xrightarrow{f_2} P_1 \xrightarrow{f_1} P_0 \xrightarrow{f_0} k \to 0,$$

where each P_i is a finitely generated graded projective, hence is a sum of shifts of A. We note that $P_0 = A$ in this case, and we write $P_i = \bigoplus_j A(-r_{ij})$ as in Corollary 24.28.

Lemma 24.32: If $0 \to V_d \to V_{d-1} \to \cdots \to V_0 \to 0$ is an exact sequence of finite-dimensional vector spaces, then $\sum_i (-1)^i \dim V_i = 0$.

Applying this lemma to the terms of degree n in the resolution as written above, we obtain the formula, valid for all $n > 0$,

$$a_n - \sum_{i=1}^{d} (-1)^{i+1} \left(\sum_j a_{n-r_{ij}} \right) = 0, \tag{3}$$

in which all r_{ij} are positive. This recursive formula, together with the initial conditions $a_n = 0$ for $n < 0$ and $a_0 = 1$, determines the Hilbert function.

Example 24.33: The q-polynomial ring $A = k_q[x, y]$, is defined by the relation $yx = qxy$. Writing operators on the right, the resolution of k is

$$0 \to A(-2) \xrightarrow{(y,-qx)} A(-1)^{\oplus 2} \xrightarrow{\binom{x}{y}} A \to k \to 0.$$

The recursive formula is $a_n = 2a_{n-1} - a_{n-2}$, and the Hilbert function is that of the commutative polynomial ring (as was clear from the start).

Example 24.34: Let $A = k\langle x, y\rangle / I$, where I is the ideal generated by the two elements $[x, [x, y]] = x^2y - 2xyx + yx^2$ and $[[x, y], y] = xy^2 - 2yxy + y^2x$. The global dimension is three, and the resolution has the form

$$0 \to A(-4) \xrightarrow{f^{(2)}} A(-3)^{\oplus 2} \xrightarrow{f^{(1)}} A(-1)^{\oplus 2} \xrightarrow{f^{(0)}} A \to k \to 0,$$

where

$$f^{(0)} = \begin{pmatrix} x \\ y \end{pmatrix}, \quad f^{(1)} = \begin{pmatrix} yx - 2xy & x^2 \\ y^2 & xy - 2yx \end{pmatrix}, \quad f^{(2)} = \begin{pmatrix} y & x \end{pmatrix}.$$

The recursive formula for the Hilbert function is $a_n - 2a_{n-1} - 2a_{n-3} + a_{n-4} = 0$.

Example 24.35: Let $A = k\langle x, y\rangle / I$, where I is the ideal generated by the element $y^2x + x^2y - x^3$. The global dimension is 2, and the resolution has the form

$$0 \to A(-3) \xrightarrow{(y^2 - x^2, x^2)} A(-1)^{\oplus 2} \xrightarrow{\binom{x}{y}} A \to k \to 0.$$

The recursive formula is $a_n - 2a_{n-1} + a_{n-3} = 0$.

Exercise 24.36: Prove that the resolutions in these past three examples are exact.

We can also describe the Hilbert series $h(t)$ conveniently in terms of the recursive formula. Because signs alternate, we can gather the terms in (3) together, to obtain a formula of the general shape

$$a_n - \sum_i a_{n-r_i} + \sum_j a_{n-s_j} = 0.$$

Let

$$q(t) = 1 - \sum_i t^{r_i} + \sum_j t^{s_j}.$$

Exercise 24.37: Prove Hilbert's theorem, that the Hilbert series of any finitely generated commutative graded ring is a rational function. Do it by writing A as a quotient of a polynomial ring P, and resolving A as a P-module.

Having expressed $h(t)$ as a rational function, we can determine the growth of the algebra. We write $q(t) = \prod(1-\lambda_i t)$, where λ_i are the reciprocal roots of $q(t)$ - the reciprocals of the roots.

Theorem 24.38: Let A be a finitely generated, connected graded algebra of finite global dimension, and let $h(t) = 1/q(t)$ be its Hilbert series.

(i) a_n has exponential growth if and only if $q(t)$ has a reciprocal root λ with $|\lambda| > 1$.

(ii) If every reciprocal root of $q(t)$ has absolute value ≤ 1, then the reciprocal roots are roots of unity, and q is a product of cyclotomic polynomials. In this case, when the reciprocal roots of q are roots of unity, then A has polynomial growth, and its GK dimension is the multiplicity of the reciprocal root 1, the order of pole of $h(t)$ at $t = 1$. Moreover, the order of pole of h at $t = 1$ is its maximal order of pole.

Proof. (i) The radius of convergence r of the rational function $h(t)$ is the minimum absolute value of its poles. So $r < 1$ if and only if $q(t)$ has a root λ of absolute value < 1.

(ii) The reciprocal roots are the nonzero roots of the polynomial $t^n q(t^{-1})$, which is a monic polynomial with integer coefficients. Furthermore, the constant term is 1/the leading coefficient of $q(t)$, which is 1, so the product of the reciprocal roots is ± 1. Since all $|\lambda_i| \leq 1$, it follows that $|\lambda_i| = 1$ for all i. It then follows from basic number theory that $q(t)$ is a product of cyclotomic polynomials. The rest of the statement is a computation; for completeness, we include the details here.

Let k denote an integer such that $\lambda_i^k = 1$ for all i, and let ζ be a primitive kth root of 1. Also, let p denote the largest multiplicity among the roots of $q(t)$. We write $h(t)$ in terms of partial fractions, say

$$h(t) = \frac{1}{q(t)} = \sum_{i,j} \frac{c_{ij}}{(1 = \zeta^i t)^j}$$

with $i = 0, \ldots, k-1$ and $j = 1, \ldots, p$, where c_{ij} are complex numbers. The binomial expansion for a negative power is

$$\frac{1}{(1-t)^j} = \sum_n \binom{n+j-1}{j-1} t^n.$$

This yields the formula

$$a_n = \sum_i c_{ij} \binom{n+j-1}{j-1} \zeta^{in},$$

where $j = 1, \ldots, p$. Thus the value of a_n cycles through k polynomial functions. For $\nu = 0, \ldots, k-1$,

$$a_n = \gamma_\nu(n) := \sum_{i,j} c_{ij} \zeta^{i\nu} \binom{n+j-1}{j-1}, \quad \text{if } n \equiv \nu \pmod{k}.$$

Because a_n takes real values at the integers $n \equiv \nu$, $\gamma_\nu(n)$ is a real polynomial. Its degree is at most $p-1$, so $GK\dim(A) \leq p$.

The coefficient of n^{p-1} in $\gamma_\nu(p)$ is

$$\frac{\gamma_{\nu p}}{(p-1)!} = \sum_i \frac{c_{ip}\zeta^{i\nu}}{(p-1)!}.$$

It is non-negative because a_n takes non-negative values. Since $h(t)$ has a pole of order p, at least one of the coefficients c_{ip} is nonzero. The coefficient vector $(\gamma_{0\,p}, \gamma_{1\,p}, \ldots, \gamma_{k-1\,p})$ is obtained from the vector $(c_{0\,p}, \ldots, c_{k-1\,p})$ by multiplying by the nonsingular matrix $(\zeta^{i\nu})$. Therefore at least one coefficient γ_{ip} is positive, and the sum $\gamma = \gamma_{0\,p} + \cdots + \gamma_{k-1\,p}$ is positive too. Since

$$\gamma = \sum_{i,\nu} c_{ip}\zeta^{i\nu} = kc_{i\,0},$$

it follows that $c_{i\,0} > 0$, which implies that h has a pole of order p at $t = 1$, and that $GK\dim(A) = p$. □

Corollary 24.39: Let $A = k \oplus \bigoplus_{i \geqslant 1} A_i$ be right (or left) Noetherian of right (or left) finite homological dimension. Then

$$h(t) = \frac{1}{q(t)},\ q(t) \in \mathbb{Z}[t]$$

where $h(t)$ is the Hilbert series and $q(t)$ is a polynomial whose roots are all roots of unity.

Proof. We must have

$$q(t) = \sum (-1)^i \dim \mathrm{Tor}_i^A(k,k) t^i$$

(i.e. the graded Euler characteristic of $\mathrm{Tor}^A(k,k)$).
By Stephenson-Zhang Theorem 24.31, A has subexponential growth (since A is assumed to be Noetherian). By Theorem 24.38, all the roots z_i of $q(t)$ must have $|z_i| \geqslant 1$; otherwise, $\sum a_n t^n$ has radius of convergence < 1 and A has exponential growth. Then the theorem part (ii) tells us that q is a product of cyclotomic polynomials, hence $q(t) \in \mathbb{Z}[t]$ is a polynomial whose roots are all roots of unity. Finally, we compute $h(t)q(t)$ and find it is 1. □

Example 24.40: In Example 24.34, $q(t) = 1 - 2t + 2t^3 - t^4 = (1-t)^3(1+t)$. All reciprocal roots here have absolute value ≤ 1, and indeed are roots of unity. The multiplicity of the root 1 is 3, so this algebra has GK dimension 3.

Example 24.41: In Example 24.35, $q(t) = 1 - 2t + t^3 = (1-t)(t^2 - t - 1)$. It has the root $\frac{-1+\sqrt{5}}{2}$, which has absolute value < 1, hence its reciprocal is > 1. By Theorem 24.38, this algebra has exponential growth, hence by Stephenson-Zhang Theorem 24.31, this algebra is not noetherian.

Conjecture 24.42 (Polishchuk-Positselski): The Hilbert series of a Koszul algebra is rational. Moreover, if both A and $A^!$ have finite GK dimension, then they have the Hilbert series of a symmetric tensor exterior.

Conjecture 24.43 (Anick): Assume A is right Noetherian. If both $\mathrm{GKdim}(A)$ and $\mathrm{hdim}(A)$ are finite, then the Hilbert series of A equals that of the symmetric algebra (in d variables y_i with degrees k_i; i.e. $(1-t^{k_1})\cdots(1-t^{k_d})$).

25 May 16 - Final class: noncommutative geometry

Recall that commutative algebra is closely related to algebraic geometry. A commutative ring R corresponds to the affine scheme $\operatorname{Spec} R$, and modules over R correspond to sheaves on $\operatorname{Spec} R$. In algebraic geometry these concepts are extended to more general non-affine schemes, while also creating powerful geometric intuition and techniques that have had a strong impact on commutative algebra.

Noncommutative geometry is an area that grew out of attempts to tie noncommutative algebra to geometry in a similar way. This has not led to as comprehensive a theory as exists in the commutative case. However, it did lead to emergence of a number of different directions, some leading to impressive results.

In this lecture we will briefly survey some of these directions. Our list is by no means complete; for example we don't discuss the direction involving tools from functional analysis (C^*-algebras) developed by A. Connes et al.

25.1 Representation varieties

Let R be a finitely generated commutative ring over $k = \bar{k}$. Then k-points of $\mathrm{Spec}(R)$ correspond to the homomorphisms of k-algebras $\mathrm{Hom}(R,k)$ i.e. to *one-dimensional* representations of the algebra R.

Note that every simple module over a finitely generated commutative algebra R is one-dimensional (Hilbert's Nullstellensatz). If R is instead a finitely generated noncommutative algebra over k, it is natural to consider the space of *all* finite-dimensional representations of R. Let us describe this space (to be denoted $\mathrm{Rep}(R)$).

First of all note that $\mathrm{Rep}(R) = \bigsqcup_{n\in\mathbb{Z}_{\geqslant 0}} \mathrm{Rep}_n(R)$, where $\mathrm{Rep}_n(R)$ is the space of n-dimensional representations of R.

Every element of $\mathrm{Rep}_n(R)$ corresponds to a homomorphism $\varphi\colon R \to \mathrm{Mat}_n(k)$ i.e. $\varphi \in \mathrm{Hom}(R, \mathrm{Mat}_n(k))$. Two homomorphisms φ_1, φ_2 define isomorphic representations iff they lie in the same orbit of GL_n, acting naturally on the space $\mathrm{Hom}(R, \mathrm{Mat}_n(k))$ (via its action on $\mathrm{Mat}_n(k)$). Since R is finitely generated, $R = k\langle x_1, \dots, x_m\rangle/I$, and

$$\mathrm{Hom}(R, \mathrm{Mat}_n(k)) \subset (\mathrm{Mat}_n)^m = k^{n^2 m} = \mathbb{A}^{n^2 m}$$

is a subset of the affine variety $(\mathrm{Mat}_n)^m$ cut out by polynomial equations. We can consider this subset as an *algebraic subvariety* of $\mathbb{A}^{n^2 m}$. We see that $\mathrm{Rep}_n(R) = \mathrm{Hom}(R, \mathrm{Mat}_n(k))/\mathrm{GL}_n(k)$ is the quotient of the algebraic variety $\mathrm{Hom}(R, \mathrm{Mat}_n(k))$ by the action of the algebraic group $\mathrm{GL}_n(k)$. Space $\mathrm{Rep}_n(R)$ is an example of an *algebraic stack*. This is a replacement for $\mathrm{Spec}(R)$.

Preprojective algebras are examples of explicit algebras with interesting representation varieties $\mathrm{Rep}\, R$.

Let Q be an oriented quiver and let $\overline{Q}$ be the corresponding double quiver. For an edge e of Q we will denote by e_+, e_- the corresponding edges of $\overline{Q}$. Let $P(Q)$ be the quiver algebra of $\overline{Q}$ modulo the relation

$$\sum_e e_- e_+ - \sum_e e_+ e_- = 0. \tag{4}$$

Remark 25.1: It's more common to quotient by the ideal generated by the elements $\theta_i := \sum_{e:i\to ?}(e_- e_+ - e_+ e_-)$ where the sum runs over all edges going out of some fixed vertex i, but this is actually equivalent to writing a single generator of the ideal: the single generator is just the sum of the θ_i, while we get each θ_i by multiplying the single generator by the idempotent of the vertex i.

For example, let Q be a cyclic quiver consisting of n vertices labeled by the elements of $\mathbb{Z}/n\mathbb{Z}$ (vertices $[i]$, $[i+1]$ are connected by the edge). Quiver $\overline{Q}$ has vertices labeled by $\mathbb{Z}/n\mathbb{Z}$, edges of this quiver are $[i] \leftarrow [i+1]$ and $[i+1] \leftarrow [i]$, $[i] \in \mathbb{Z}/n\mathbb{Z}$.

Pick $\zeta \in k$ of order n and consider the action $\mathbb{Z}/n\mathbb{Z} \curvearrowright k[x,y]$ given by $[1]\cdot x = \zeta x$, $[1]\cdot y = \zeta^{-1} y$. We have $(\mathbb{Z}/n\mathbb{Z})\# k[x,y] \xrightarrow{\sim} P(Q)$, the isomorphism is given by:

$$1\otimes x \mapsto \sum_{[i]\in\mathbb{Z}/n\mathbb{Z}} e_{[i+1]\leftarrow[i]},\ 1\otimes y \mapsto \sum_{[i]\in\mathbb{Z}/n\mathbb{Z}} e_{[i]\leftarrow[i+1]},\ [1]\otimes 1 \mapsto \sum_{[i]\in\mathbb{Z}/n\mathbb{Z}} \zeta^i e_{[i]}.$$

The isomorphism above induces the equivalence between the categories of $P(Q)$ and $(\mathbb{Z}/n\mathbb{Z})\# k[x,y]$-modules. Let us describe this equivalence explicitly. Module $(M_{[i]})_{[i]\in\mathbb{Z}/n\mathbb{Z}}$ over $P(Q)$ goes to $M := \bigoplus_{[i]\in\mathbb{Z}/n\mathbb{Z}} M_{[i]}$, where the action of $[1] \in \mathbb{Z}/n\mathbb{Z}$ on $M_{[i]}$ is given by ζ^i and the action of $x\colon M_{[i]} \to M_{[i+1]}$ is given by $e_{[i+1]\leftarrow[i]}$, the action of $y\colon M_{[i+1]} \to M_{[i]}$ is given by $e_{[i]\leftarrow[i+1]}$. The condition $\sum_{[i]\in\mathbb{Z}/n\mathbb{Z}} e_{[i]\leftarrow[i+1]} e_{[i+1]\leftarrow[i]} = \sum_{[i]\in\mathbb{Z}/n\mathbb{Z}} e_{[i+1]\leftarrow[i]} e_{[i]\leftarrow[i+1]}$ precisely corresponds to the fact that x and y commute.

This example can be generalized as follows. Recall that finite subgroups Γ in $\mathrm{SL}(2,k)$ correspond to simply laced Dynkin graphs D (this is known as McKay correspondence, see [15]). Let $\widehat{D}$ be the affine Dynkin graph; the vertices of $\widehat{D}$ are in bijection with irreps of Γ (see [15]). Then (see [7])

$$P(\widehat{D}) \sim \Gamma \# k[x,y]$$

where the $\sim$ is Morita equivalence. It sends a $\Gamma\# k[x,y]$-module M to $\bigoplus_v M_v$, where $M_v = [M : \rho_v] = \mathrm{Hom}_\Gamma(\rho_v, M)$ and ρ_v is the irreducible representation of Γ corresponding to the vertex v.

Remark 25.2: Note that the algebras $P(\widehat{D})$, $\Gamma\# k[x,y]$ are not isomorphic in general (they are isomorphic for $\Gamma = \mathbb{Z}/n\mathbb{Z}$).

Let us describe the representation variety of the algebra $R = P(Q)$. Let us first of all recall that $P(Q)$ is a certain quotient of the path algebra of the quiver $\overline{Q}$ so every representation of $P(Q)$ can be considered as a representation of the quiver $\overline{Q}$ such that (4) holds. So, $\mathrm{Rep}(P(Q)) = \bigsqcup_{d_v \in \mathbb{Z}_{\geqslant 0}} \mathrm{Rep}_{d_v}(P(Q))$, where $\mathrm{Rep}_{d_v}(P(Q))$ is the space of representations (M_v) of $\overline{Q}$ such that (4) holds and $\dim M_v = d_v$ (considered up to an isomorphism). Explicitly, let

$$\widetilde{\mathrm{Rep}}_{d_v}(P(Q)) \subset \prod_{e\colon v\to v'} \mathrm{Mat}_{d_v,d_{v'}} \times \mathrm{Mat}_{d_{v'},d_v} \tag{5}$$

be the subvariety consisting of collections of maps such that (4) holds. Then

$$\mathrm{Rep}_{d_v}(P(Q)) = \widetilde{\mathrm{Rep}}_{d_v}(P(Q))/\prod_v \mathrm{GL}(d_v).$$

Note now that the RHS of (5) is a *symplectic vector space* (we identify $\mathrm{Mat}_{d_v,d_{v'}} \times \mathrm{Mat}_{d_{v'},d_v}$ with $T^* \mathrm{Mat}_{d_v,d_{v'}}$ via the trace form), and $G = \prod \mathrm{GL}(d_v)$ acts on it preserving the symplectic structure; the equations (4) are zeroes of the moment map $\mu\colon \prod_{e\colon v\to v'} T^* \mathrm{Mat}_{d_v,d_{v'}} \to \prod_v \mathrm{Mat}_{d_v}$ for the $\prod_v \mathrm{GL}_{d_v}$ action. So, $\mathrm{Rep}_{d_v}(P(Q))$ is obtained from the RHS by Hamiltonian reduction i.e.

$$\mathrm{Rep}_{d_v}(P(Q)) = \mu^{-1}(0)/\prod_v \mathrm{GL}_{d_v}$$

is the *Hamiltonian reduction* of the symplectic vector space $\prod_{e\colon v\to v'} T^* \mathrm{Mat}_{d_v,d_{v'}}$ by $\prod_v \mathrm{GL}_{d_v}$.

Remark 25.3:
Ginzburg introduced a notion of a *quiver with potential* (see [8, Section 4.2]). Potential is an *element* of the vector subspace F_{cyc} of the quiver algebra $F = kQ$ of a quiver Q generated by *cyclic* paths.
For any edge $e \in D$ there exists a map $\frac{\partial}{\partial e}\colon F_{\mathrm{cyc}} \to F$ defined as follows: given a cyclic path $\Phi = e_{i_1} e_{i_2} \ldots e_{i_r}$, we put

$$\frac{\partial \Phi}{\partial e} := \sum_{\{s \mid e_{i_s} = e\}} e_{i_{s+1}} e_{i_{s+2}} \ldots e_{i_r} e_{i_1} e_{i_2} \ldots e_{i_{s-1}}.$$

One can then consider the quotient algebra $\mathfrak{A}(Q,\Phi) := F/(\frac{\partial\Phi}{\partial e})_{e\in Q}$, where $(\frac{\partial\Phi}{\partial e})_{e\in D}$ is the two-sided ideal generated by the elements $\frac{\partial\Phi}{\partial e} \in F$ (see [8, Equation (4.2.1)]).
The variety $\widetilde{\mathrm{Rep}}_{d_v} \mathfrak{A}(Q,\Phi)$ can be described as follows. Consider the space $\widetilde{\mathrm{Rep}}_{d_v}(F)$ (that is just a vector space), element Φ defines a map $\widehat{\Phi}\colon \widetilde{\mathrm{Rep}}_{d_v}(F) \to \prod_v \mathrm{Mat}_{d_v}$ sending a representation ρ to $\rho(\Phi)$ (recall that $\Phi \in F_{\mathrm{cyc}}$). We obtain a functional $\phi := \mathrm{tr}\, \widehat{\Phi}\colon \widetilde{\mathrm{Rep}}_{d_v}(F) \to \mathbb{C}$. One can show that the variety $\widetilde{\mathrm{Rep}}_{d_v} \mathfrak{A}(Q,\Phi)$ is the *critical locus* of $\mathrm{tr}\, \widehat{\Phi}$ (cf. [8, Section 2.3]).
Let us describe how to obtain preprojective algebras via algebras with potentials (see [8, Example 4.3.3]). Consider the quiver $\overline{Q}^{\mathrm{loop}}$ obtained from $\overline{Q}$ by attaching an additional edge loop, t_v, for every vertex v. We can identify the quiver algebra $k\overline{Q}^{\mathrm{loop}}$ with $k\overline{Q} * k[t]$, where $*$ corresponds to the free product of associative k-algebras (this isomorphism sends t to $\sum_v t_v$). Consider the potential

$$\Phi := \sum_v t_v \cdot \sum_e (e_+ e_- - e_- e_+) = t \sum_e (e_+ e_- - e_- e_+).$$

We then have $\mathfrak{A}(k\overline{Q}^{\mathrm{loop}}, \Phi) = P(Q)[t]$ (see [8, Equation (4.3.4)]), so $\widetilde{\mathrm{Rep}}_{d_v} \mathfrak{A}(k\overline{Q}^{\mathrm{loop}}, \Phi) = \widetilde{Rep}_{d_v}(P(Q)) \times \prod_v \mathrm{Mat}_{d_v}$. The potential ϕ that we construct above is given by the formula:

$$\prod_{e\colon v\to v'} T^* \mathrm{Mat}_{d_v,d_{v'}} \times \prod_v \mathrm{Mat}_{d_v} \ni (v,\xi) \mapsto \mathrm{tr}(\mu(v)\xi) \in \mathbb{C}$$

in this case.

25.2 Weyl algebra and deformations

Many interesting algebras have no nonzero finite-dimensional representations (in particular, $\mathrm{Rep}\, R = \{0\}$ is not interesting in this case). For example, the Weyl algebra $W = \mathbb{C}\langle x,y\rangle/\langle xy - yx - 1\rangle$ doesn't. You can see this by noting that $\mathrm{tr}(xy - yx) = 0$ (on every finite-dimensional representation), while $\mathrm{tr}(1_n) = n$.

We can study noncommutative geometry by deforming from the commutative case, this procedure is also called deformation quantization. Consider

$$W_\hbar = \mathbb{C}[\hbar]\langle x, y\rangle / \langle xy - yx - \hbar\rangle.$$

When we take $\hbar = 1$, we recover W, and when we take $\hbar = 0$, we get $\mathbb{C}[x, y]$, which is commutative.

Other examples of deformations:

If X is a smooth affine algebraic variety over a field k, we can consider $\mathrm{Diff}_\hbar(X)$, the asymptotic differential operators. This is W when $X = \mathbb{A}^1$. If $\hbar = 0$, we get $\mathcal{O}(T^*X) = \mathrm{Sym}_{\mathcal{O}(X)}(\mathrm{Der}(\mathcal{O}(X)))$.

If $\mathfrak{g}$ is a Lie algebra, let $U_\hbar(\mathfrak{g}) = k\langle\mathfrak{g}\rangle / xy - yx = \hbar x$. If $\mathfrak{g} = \mathfrak{gl}_n$, then define

$$\bar{U_\hbar}(\mathfrak{g}) = U_\hbar \otimes_{Z(U_\hbar(\mathfrak{g}))} k[\hbar]$$

so

$$\bar{U_0} = \mathcal{O}(\mathcal{N})$$

where $\mathcal{N}$ is the nilpotent matrices.

There's also the spherical rational double affine Hecke algebra (DAHA) $A_\hbar$, also called rational Cherednik algebra, where $A_0 = \mathcal{O}((\mathbb{A}^2)^n / S_n)$.

How can we define deformations for non-affine varieties? Previously we deformed the algebra of functions on X, now we need to deform the structure sheaf of X. There is no obvious way to make a deformation into a sheaf. Here are some ways:

a) We can work with a formal parameter. If A is flat over $k[[\hbar]]$ and complete in the $\hbar$-topology, let $A_0 = A/\hbar$, which is commutative. Exercise: a) if $\bar{a} = a \pmod{\hbar}$ is invertible, then so is a. b) If $U \subset \mathrm{Spec}(A)$ is open, then $\{a \in A | \bar{a}|_U \text{ is invertible}\}$ is a localizing class. c) Localizations form a Zariski sheaf on $\mathrm{Spec}(A)$.

b) $\mathrm{Diff}(X)$ can be made into a sheaf on T^*X with conical topology ($U \subset \mathrm{Spec}(T^*X)$ is open in a conical topology if it is open in Zariski topology and invariant under dilation).

c) In characteristic p, for $W = k\langle x, y\rangle / xy - yx = 1$ has $x^p, y^p \in Z(W)$. Hence, W is a sheaf on $\mathrm{Spec}\, k[x^p, y^p]$.

25.3 $\mathrm{Coh}(X)$ and $D^\flat(\mathrm{Coh}(X))$

The previous two subsections described approaches closely tied to the usual commutative algebraic geometry. The next relies on it as a source of motivation for conjectures rather than trying to relate a noncommutative structure to a specific commutative ring or variety.

An algebraic variety X can be studied via the category $\mathrm{Coh}(X)$ on $D^\flat(\mathrm{Coh}(X))$.

If $X = \mathrm{Spec}(R)$ is affine, then $\mathrm{Coh}(X) = R$-mod. If Y is projective over k and $Y = \mathrm{Proj}(A)$, $A = \bigoplus_{n\geq 0} A_n$, $A_0 = k$, then $\mathrm{Coh}(Y)$ is the Serre quotient of graded finitely generated A-modules by graded finite-dimensional A-modules.

Theorem 25.4 (Serre): Let X be an algebraic variety over a field k. Then X is smooth iff $\mathrm{Coh}(X)$ has finite homological dimension, i.e. $\mathrm{Ext}^n(F, G) = 0$ for all $n > d$ and all $F, G \in \mathrm{Coh}(X)$.

It is also known that X is projective iff $\mathrm{Ext}^n(F, G)$ is finite-dimensional for all n, $F, G \in \mathrm{Coh}(X)$.

Let X be a smooth affine variety. Then Ω^i_X, the i-forms, is $\mathrm{HH}_i(\mathcal{O}(X))$. Recall that HH^i and HH_i are Morita invariant. They can also be defined starting from a category: $\mathrm{HH}^i = \mathrm{Ext}^i(\mathrm{Id}, \mathrm{Id})$ where Id is the identity functor, while HH_i relates to the tensor of bimodules.

For X smooth and projective,

$$\mathrm{HH}_i(\mathrm{Coh}(X)) = \bigoplus_{q-p=i} H^p(X, \Omega^q) \simeq H^i_{\mathrm{dR}}(X)$$

where the last equivalence is from the Hodge theorem, and the first one is known as the Hochschild-Kostant-Rosenberg isomorphism.

To recover H^*_{dR} for nonprojective X, we can use cyclic homology. The bar complex for HH_* has cyclic symmetry:

$$C : a_0 \otimes \cdots a_n \to (-1)^n a_n \otimes a_0 \cdots \otimes a_{n-1}.$$

Then $\mathrm{Bar}/(C - \mathrm{Id})\mathrm{Bar}$ inherits the differential from the bar complex. Its cohomology is

$$\mathrm{HC}_n(A) = \Omega^n / d\Omega^{n-1} \oplus \bigoplus_{i \geq 1} H^{n-2i}_{\mathrm{dR}}(X)$$

and

$$\mathrm{HC}^{\mathrm{per}}_n = \varinjlim \mathrm{HC}_{n+2i} = \bigoplus_{i=-\infty}^{\infty} H^{n+2i}_{\mathrm{dR}}(X).$$

For a smooth projective dimension n variety X we have Serre duality:

$$\mathrm{Ext}^i(F, G) \simeq \mathrm{Ext}^{n-i}(G, F \otimes K_X)^*.$$

Definition 25.5 (Bondal-Kapranov): Let C be a finite type k-linear triangulated category. (Finite type means $\dim_k \mathrm{Ext}^*(A, B) < \infty \, \forall A, B$.) For example, $C = D^b(A\text{-mod})$ where A is finite-dimensional and of finite homological dimension.
A Serre functor is a functor $S : C \to C$ and an isomorphism $\mathrm{Hom}(A, B) \simeq \mathrm{Hom}(B, S(A))^*$. The Yoneda lemma implies that if S exists, it is unique.

Example 25.6: For $C = D^b(\mathrm{Coh}(X))$, $S : F \mapsto F \otimes K_x[n]$ is a Serre functor.

The Hodge theorem can be restated as a claim that a spectral sequence $\mathrm{HH}_*(\mathrm{Coh}(X)) \implies \mathrm{HC}^{\mathrm{per}}_*(\mathrm{Coh}(X))$ degenerates for smooth projective X. The following striking generalization was proposed (in a slightly different form) by Kontsevich and Soibelman (see [11]).

Conjecture 25.7: The above spectral sequence degenerates for any dg-category of finite type over k.

This was partly proved by Dmitry Kaledin and Akhil Mathew, see [10], [14].

25.4 Artin-Schelter regular algebras and noncommutative projective geometry

A projective variety $X \subset \mathbb{P}^n_k$ is determined by its homogeneous coordinate ring A, a graded commutative ring such that $X = Proj(A)$. An important invariant of X is the abelian category $Coh(X)$ of coherent sheaves on X. It can be realized as a Serre quotient $A - mod^{gr}_{fg} / A - mod^{gr}_{fd}$ where $A - mod^{gr}_{fg}$ is the category of finitely generated graded A-modules and $A - mod^{gr}_{fd}$ is the Serre subcategory of finite dimensional graded modules.

One can study noncommutative graded ring that share basic features with commutative ones, thinking of them as homogeneous coordinate rings of (yet to be defined) noncommutative projective varieties. Some beautiful results in that direction were obtained by Artin, Schelter and others in 1990's.

So consider a nonnegatively graded algebras over a field $A = \oplus A_n$, $A_0 = k$. Assuming A is Noetherian, the category $A - mod^{gr}_{fg}$ is abelian, so one can consider $A - mod^{gr}_{fg} / A - mod^{gr}_{fd}$, the category of coherent sheaves on the purported noncommutative $Proj$ of A.

One defines a *point module* over A as a cyclic graded module with Poincare series $1/(1-t)$. In the case when A is commutative and generated by A_1 point modules are in bijection with points of $X = Proj(A)$.

Several important classification results are achieved by considering point modules and a usual (commutative) algebraic variety arising as the moduli space of point modules.

We briefly mention a sample classification problem. Recall that a commutative A as above has finite homological dimension (equivalently, is regular) iff it is a polynomial algebra. Assuming it is generated by A_1, we get the homogeneous coordinate ring of $X = \mathbb{P}^n_k$. Generalization of this simplest projective variety leads to the following definition.

An algebra A as above is called *Artin-Schelter regular* if it has finite homological dimension d, a finite GK dimension and $Ext^i_A(k, A) = 0$ for $i \neq d$, while $Ext^d_A(k, A)$ is one dimensional.

The work of Artin-Schelter and Artin-Tate-van den Bergh achieved classification of AS regular algebras of dimensions two and three (noncommutative projective lines and planes). We will not present the answer, but mention that it involves beautiful and rather elementary algebro-geometric data, such as an elliptic curve with an automorphism (see [17] and references therein), arising in the process of classification of point modules over A.

References

[1] Michael Artin. Noncommutative Rings class notes. 1999.

[2] Asher Auel, Eric Brussel, Skip Garibaldi, and Uzi Vishne. Open problems on central simple algebras. *Transformation Groups*, 16(1):219–264, 2011.

[3] J. Backelin. *On the rate of growth of the homologies of Veronese subrings.* in Algebra, Algebraic Topology and their Interactions. J.-E. Roos ed., Lect. Notes Math., 1183, pp. 79–100, Springer-Verlag, Berlin/New York, 1986.

[4] A.A. Beilinson, V.A. Ginzburg, and V.V. Schechtman. Koszul duality. *Journal of Geometry and Physics*, 5(3):317–350, 1988.

[5] Alexander Beilinson, Victor Ginzburg, and Wolfgang Soergel. Koszul Duality Patterns in Representation Theory. *Journal of the American Mathematical Society*, 9(2):473–527, 1996.

[6] R. Bezrukavnikov. *Koszul property and Frobenius splitting of Schubert varieties.* arXiv preprint alg-geom/9502021, 1995.

[7] William Crawley-Boevey, Martin P. Holland. *Noncommutative deformations of Kleinian singularities.* Duke Math. J. 92 (1998), no.3, 605–635.

[8] Victor Ginzburg. *Calabi-Yau algebras.* arXiv:math/0612139.

[9] Alexander Grothendieck. *Sur quelques points d'algèbre homologique.* Tohoku Mathematical Journal, 1957.

[10] D. Kaledin. *Spectral sequences for cyclic homology.* In Algebra, geometry, and physics in the 21st century, volume 324 of Progr. Math., pages 99–129. Birkhäuser/Springer, Cham, 2017.

[11] Maxim Kontsevich and Yan Soibelman. *Notes on A-infinity algebras, A-infinity categories and non-commutative geometry. I.* arXiv preprint math/0606241 (2006).

[12] Günter R. Krause and Thomas H. Lenagan. *Growth of Algebras and Gelfand-Kirillov Dimension.* Number 22 in Graduate Studies in Mathematics. American Mathematical Society, 2000.

[13] Saunders Mac Lane. *Categories for the working mathematician.* Number 5 in Graduate texts in mathematics. Springer-Verlag, 1971.

[14] Akhil Mathew. *Kaledin's degeneration theorem and topological Hochschild homology.* Geometry Topology 24.6 (2020): 2675-2708.

[15] John McKay. *Graphs, singularities, and finite groups.* Proc. Symp. Pure Math. Vol. 37. No. 183. 1980

[16] Inamdar, S. P., Mehta, V. B. (1994). Frobenius splitting of Schubert varieties and linear syzygies. American Journal of Mathematics, 116 (6), 1569-1586.

[17] Daniel Rogalski, *Artin-Schelter Regular Algebras.* arXiv:2307.03174, 2023.

[18] Charles A. Weibel. *An introduction to homological algebra.* Number 38 in Cambridge studies in advanced mathematics. Cambridge Univ. Press, Reprint. 1997, transf. to digital print edition, 2003.