ALGEBRA QUAL PREP: LINEAR ALGEBRA SOLUTIONS

TONY FENG

These are hints/solutions/commentary on the problems. They are not a model for what to actually write on the quals.

1. 2014 Fall, Afternoon #4

(a) Rational canonical form says that T is a direct sum of "companion matrices", which act like multiplication by T on $k[T]/f(T)$. If $f(T)=g(T)h(T)$ with g and h coprime, then

$$k[T]/f(t) \simeq k[T]/g(T) \oplus k[T]/h(T).$$

Therefore, it suffices to show:

- A (finite) direct sum of operators is semisimple if and only if each factor is.
- If f is a power of an irreducible, then multiplication by T on $k[T]/f(T)$ is semisimple if and only if f is irreducible.

First we consider the "only if" direction. By induction we restrict our attention to $V = V_1 \oplus V_2$, a T-invariant direct sum, and we need to show that if $W \subset V_1$ is T-invariant then it has a complement $W^{\perp} \subset V_1$. By assumption we can take a complement U for W in V. We then need to produce a subspace of V_1. We should either take the projection of U in V_1 or the intersection of U with V_1, and the content of the problem is to decide which is correct.

By definition, any $v \in V$ can be uniquely written as $v = w + u$ for $w \in W$ and $u \in U$. Also by definition, if $v \in V_1$ then $u \in V_1$. Hence any $v \in V$ can be uniquely written as $w + u$ for $w \in W$ and $u \in V_1 \cap U$. So we conclude that $(V_1 \cap U)$ is a complement for W in V_1.

Next we consider the "if" direction. We are reduced to the case of two summands by induction, say $V = V' \oplus V''$ with each factor T-stable. Consider the sequence

$$0 \to V'' \to V \to V' \to 0.$$

If $W \subset V$ is T-stable, then its quotient in V' is T-stable, hence admits a complement W'. Its kernel is also T-stable, hence admits a complement W''. Then check that $W' \oplus W''$ is a complement for W.

LPT 1. It would *not* have been good to consider $W \cap V'$ and $W \cap V''$. (These both have T-stable complements, but the sum of the complements is not a complement for W. Exercise: find an example.) As a general principle, it's better to work with filtrations than summands, at least when intersecting.

Now we are reduced to the case $f(T) = g(T)^k$ for some irreducible $g(T)$. If $k = 1$, then we claim that $k[T]/f(T)$ has no non-trivial T-stable subspaces. If it did, the characteristic polynomial of T on that subspace would be a polynomial strictly dividing $f(T)$, so it would be 1, so that subspace would be 0. Conversely, if $k > 1$, then filter the space as

$$0 \to \text{Image}(\times g(T)^{k-1}) \to k[T]/g(T)^k \to \text{coker}(\times g(T)^{k-1})) \to 0.$$

Suppose Image$(\times g(T)^{k-1})$ had a T-stable complement. That complement would project isomorphically onto the cokernel, which is necessary isomorphic to $k[T]/g(T)$ as a vector space with action of T. But then $g(T)^{k-1}$ would annihilate the whole vector space, a contradiction.

(b) Jordan cannoical form says that T is conjugate to a direct sum of *Jordan blocks*, which have the form

$$\begin{pmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \ddots & \ddots \\ & & & \lambda \end{pmatrix}$$

Take S to be the diagonal λ, and N to be the rest.

(c) Suppose $S + N = S' + N'$.

LPT 2. It may be tempting to write $S - S' = N' - N$, and claim that the left side is semisimple and the right side is nilpotent. This argument is **wrong**. A sum of nilpotent operators isn't necessarily nilpotent, and a sum of semisimple operators isn't necessarily semisimple. (Exercise: find examples!) It's **extremely** important to remember that these sorts of statements are true only when you know that the operators *commute*. (Exercise: prove it in that case!)

Since S' commutes with N', it commutes with $S' + N' = T$. Therefore, it preserves the decomposition of T into Jordan blocks, i.e. it preserves the Jordan blocks. Now comes the key point: *restricted to a Jordan block S* commutes with everything, since we chose it to be scalar there. So S' commutes with S, and also N. Since T commutes with N, so does N'. *Now* we are justified in saying that $S - S'$ is semisimple and $N' - N$ is nilpotent, so their equality forces both to be 0.

2. 2010 Spring, Morning # 2

LPT 3. Try out some easier cases to get intuition. (Replace 6 and 9 with smaller numbers.)

LPT 4. item It's useful to think of a basis for $V \otimes W$ as a "box", with bases for V and W along the two axes.

Every aspect of this problem is captured by the diagram

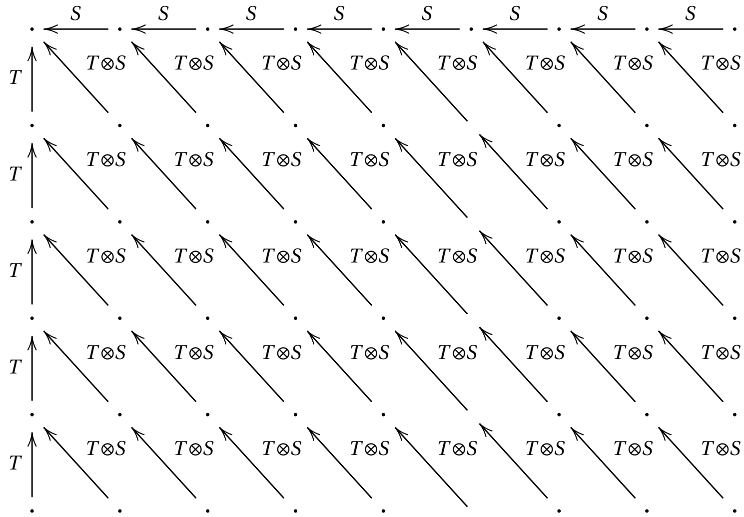

which gives a pictorial representation of a basis for $V \otimes W$. Exercise: think about the diagram until you see why.

(i) 6
(ii) 14
(iii) See diagram.

3. 2011 Spring, Morning # 3

(a) There are $\frac{q^2-q}{2}$ monic irreducible quadratics, and $\frac{q^3-q}{3}$ monic irreducible cubics. The reason is that there is a map from $\mathbf{F}_{q^2} - \mathbf{F}_q$ to the set of monic irreducible quadratics given by "minimal polynomial", which is 2:1. Similarly for the cubics.

(b) We don't need (a). By rational canonical form, we have to specify either:
- Three monic linear polynomials, which are equal, excepting x because it corresponds to a non-invertible matrix. ($q-1$ possibilities)
- A monic quadratic and a monic linear dividing it, all with non-zero constant term. ($(q-1)^2$ possibilities)
- A monic cubic polynomial with non-zero constant term. (q^3-q^2 possibilities)

Total: q^3-q.

4. 2011 Spring, Afternoon # 1

(a) The matrix is conjugate to a unique matrix in *Jordan canonical form*, meaning it is a sum of blocks of the form

$$\begin{pmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \ddots & \ddots \\ & & & \lambda \end{pmatrix}$$

Proof: view the vector space as a module over $k[T]$. Since it's finitely generated, it decomposes as

$$\bigoplus k[T]/f_i(T).$$

Using the Chinese remainder theorem, if $f(T) = g(T)h(T)$ then

$$k[T]/f(t) \simeq k[T]/g(T) \oplus k[T]/h(T).$$

So we can assume that each $f_i(T)$ has only one prime factor. Since we are over an algebraically closed field, this means that

$$f_i(T) = (T - \lambda_i)^n.$$

Making the change of variables $T' = T - \lambda_i$, we find that $k[T]/f_i(T) \simeq k[T']/(T')^n$, which for the natural basis being powers of T' has matrix

$$\begin{pmatrix} 0 & 1 & & \\ & 0 & 1 & \\ & & \ddots & \ddots \\ & & & 0 \end{pmatrix}$$

Then add the λ_i back in to finish.

(b) (b) Clear from (a).

5. 2012 Spring Afternoon # 7

By rational canonical form, M is the direct sum of operators isomorphic to "mutliplication by T" on $k[T]/f(T)$. Let's call this an "elementary block". Since the transpose of a direct sum is the direct sum of the tranposes, it suffices to show that each elementary block is isomorphic to its transpose.

Elementary blocks are *characterized* by their minimal polynomials, i.e. the elementary block $k[T]/f(T)$ is classified by its minimal polynomial $f(T)$. The conjugate matrix can be thought of as a matrix for the *dual* to "multiplication by T on $k[T]/f(T)$". If we can show that any linear transformation and its dual have the *same* minimal polynomial, then this simultaneously tells us that:

(1) the dual linear transformation is associated to a single elementary block (because its minimal polynomial has the same degree as the dimension),
(2) the dual linear transformation is associated to $k[T]/f(T)$.

So it suffices to establish that any linear transformation T and its dual T^* have the same minimal polynomial. Since dualization commutes with addition and (anti)commutes with multiplication, $f(T)^* = f(T^*)$. This makes it clear that if $f(T) = 0$ then $f(T^*) = 0$. Since $T = (T^*)^*$, that gives the other direction for free, and we are done.

Remark 5. I recommend thinking about the analogy between these argument and one which is probably more familiar to you: any finite abelian group is isomorphic to its (Pontrjagin) dual.

6. 2014 Spring Afternoon # 5

Already done.

7. 2010 Fall, Afternoon # 5

As the hint suggests, you want to prove that there is a basis with "intersection matrix"

$$\begin{pmatrix} & & -1 & \\ & & & -1 \\ 1 & & & \\ & 1 & & \end{pmatrix}$$

LPT 6. This fact generalizes to a non-degenerate quadratic form in n dimensions. It is the most fundamental fact to know about symplectic linear algebra. Make sure you know how to prove it in general.

The only thing I memorize about the proof of this fact is that "the greedy algorithm works". This means that if you just do the most naïve thing at each step, then you'll succeed.

(1) First, pick any u_1.
(2) Then pick any u_2 such that $\langle u_1, u_2\rangle = 0$. (Why is this possible?)
(3) Next, pick any v_1 such that $\langle u_1, v_1\rangle = 1$ and $\langle u_2, v_1\rangle = 0$. (Why is this possible?)
(4) Finally, pick v_2 such that $\langle u_2, v_2\rangle = 1$. You might not have $\langle u_1, v_2\rangle = 0$, but then you can just subtract off v_1 from v_2.

(a) Pick a basis (u_1, u_2, v_1, v_2) for U with the above intersection matrix. Pick a basis (u_1', u_2', v_1', v_2') for U_0 with the above intersection matrix. A linear transformation can be defined on a basis, so define g by sending $u_i \mapsto u_i'$ and $v_i \mapsto v_i'$. Why is this $g \in G$?
(b) Follow the hint. First count permissible (u_1, u_2):
   (a) There are $q^4 - 1$ choices for u_1.
   (b) Then there are $q^3 - q$ choices for u_2.

   So we get $(q^4-1)(q^3-1)$ pairs. This counts isotropic subspaces *with a basis*. To get rid of the basis, divide by $GL_2(\mathbf{F}_q)$, which has size $(q^2-1)(q^2-q)$. (If you don't know why, prove it!)

8. 2013 Fall, Morning # 5

LPT 7. It is often useful to think of a bilinear form

$$B: V \times V \to k$$

as instead a linear map

$$B(v,-): V \to V^*$$

sending $v \mapsto B(v,-)$. As an exercise, you should check that non-degeneracy of B is equivalent to $B(v,-)$ being an isomorphism (for finite-dimensional V).

(a) The map $V \xrightarrow{\sim} V^* \twoheadrightarrow U^*$ has kernel $U^\perp$.
(b) We proceed by induction on n. Pick a basis $w_1, \dots, w_n$ for W. We claim that we can choose w_1' such that $w_1' \perp \langle w_2, \dots w_n\rangle$ and $\langle w_1, w_1'\rangle = 1$. If we can do this, then $\langle w_1, w_1'\rangle^\perp \supset \langle w_2, \dots, w_n\rangle$ and still has a non-degenerate symplectic form (check it!).

   To check that the claimed w_1' exists, observe that since $\dim\langle w_2, \dots w_n\rangle^\perp = n+1$, there exists such w_1' with $\langle w_1, w_1'\rangle \neq 0$, and by rescaling we can make it 1.

(c) Simple generalization of 2010 Fall, Afternoon # 5.

9. 2013 Spring, Morning # 5

(i) Since $W^\perp \supset W$ and $\dim W^\perp = 2n - \dim W$, if $\dim W \leq n$ then W is certainly maximal. Conversely, check that the symplectic form on W descends to a non-degenerate symplectic form on $W^\perp/W$. If this is non-zero, then the lift of any non-zero vector can be added to W to produce a larger isotropic space.

(ii) Follow the hint (we've discussed how to prove it - "the greedy algorithm works"). I did not see a nice way to compute that $\det g = 1$, but I think there should be one. Here is a way by brute force. We conclude that g is upper-triangular, say

$$g \sim \begin{pmatrix} A & B \\ 0 & D \end{pmatrix}.$$

Then use the condition

$$gJg^{-1} = J \qquad J \sim \begin{pmatrix} 0 & -\mathrm{Id} \\ \mathrm{Id} & 0 \end{pmatrix}.$$

This leads to $\det B = \det D = \det A$, so $\det G = (\det A)^2 = 1$.

(iii) If it did, then consider gW with $\det g \neq 1$. By assumption, $gW = hW$ for some h with $\det h = 1$. Then $h^{-1}g$ preserves W, yet does not have determinant 1.

10. 2014 Spring, Morning # 5

(i) We have two isomorphisms

$$\omega_1(-, v) \colon V \xrightarrow{\sim} V^*$$

and

$$\omega_2(-, v) \colon V \xrightarrow{\sim} V^*.$$

Then A is the operator fitting into the diagram

$$\begin{array}{ccc} V & \xrightarrow{\omega_2(-,v)} & V^* \\ {\scriptstyle \alpha}\big\downarrow & & \big\| \\ V & \xrightarrow[\omega_1(-,v)]{} & V^* \end{array}$$

For the second part,

$$\omega_1(Av, w) = -\omega_1(w, Av) = -\omega_2(w, v) = \omega_2(v, w) = \omega_1(v, Aw).$$

(ii) The problem becomes clear after doing some small examples:

$$\omega_1(v, Av) = \omega_2(v, v) = 0$$
$$\omega_1(v, A^2 v) = \omega_1(Av, Av) = 0.$$

etc.

(iii) For $v_\lambda \in V_\lambda$ and $v_\mu \in V_\mu$, argue that

$$\omega_1(v_\lambda, v_\mu) = 0$$

by arguing on the power of $(A-\lambda)$ killing v_λ and $(A-\mu)$ killing v_μ. Note that

$$\omega_1((A-\lambda)^n v_\lambda, v_\mu) = \omega_1(v_\lambda, (A-\lambda)^n v_\mu).$$

For large n, the left side is 0, so the right side is 0 as well. By writing

$$(A-\lambda)^n = (A-\mu+\mu-\lambda)^n$$

we see that $(A-\lambda)^n$ is *invertible* on V_μ. So the fact $\omega_1(v_\lambda, (A-\lambda)^n v_\mu) = 0$ for all v_μ implies that $\omega_1(v_\lambda, v_\mu)$ for all v_μ.

(iv) By (iii), we can reduce to the case where A has a single generalized eigenspace. The minimal polynomial $(T-\lambda)^n$ where n is the length of the longest Jordan block. By (ii), the Jordan block is isotropic so can have only half the dimension.