QUALIFYING EXAM SYLLABUS

GABRIEL BEINER

Date and Time. Friday 15 November 2024, 1 PM–4 PM.

Location. 961 Evans Hall.

Committee. Ian Agol, Michael Hutchings (advisor), Sung-Jin Oh, and Fraydoun Rezakhanlou.

1. Major Topic: Symplectic Geometry (Geometry/Topology)

- **Symplectic Linear Algebra:** symplectic vector spaces; symplectic groups; Maslov index; compatible triples. ([McS1] Ch. 2)
- **Symplectic Manifolds:** Moser's trick, Moser isotopy, and Darboux's theorem; symplectic, co/isotropic, and Lagrangian submanifolds; Weinstein neighbourhood theorems; basic examples. ([Can] Ch. 3, 7–9, [McS1] Ch. 3)
- **Classical Mechanics:** Hamiltonian vector fields; Hamilton's equations; Poisson manifolds; integrable systems; Arnold–Liouville theorem. ([Can] Ch. 18)
- **Contact Manifolds:** contact structures, contact forms, and Reeb vector fields; symplectization; Liouville vector fields, contact type boundary, Liouville domains, and symplectic completion. ([Can] Ch. 10–11, [McS1] §3.5)
- **Examples:** Weinstein and Stein domains; Kähler manifolds; symplectic blowups and blowdowns; fibre connected sums. ([McS1] Ch. 7)
- **Hamiltonian Actions:** moment maps; symplectic reduction; Atiyah–Guillemin–Sternberg convexity theorem (sketch), toric manifolds and the Delzant construction; Duistermaat–Heckman formula. ([McS1] Ch. 5)
- **J-holomorphic Curves:** unique continuation; simple and somewhere injective curves; intersection positivity; adjunction inequality; moduli space of J-holomorphic curves, transversality (sketch), Gromov compactness and bubbling (sketch); Gromov nonsqueezing. ([McS2] Ch. 2–4)

References. Cannas da Silva, *Lectures on Symplectic Geometry* [Can], McDuff and Salamon, *Introduction to Symplectic Topology* (3rd ed.) [McS1], McDuff and Salamon, *J-holomorphic Curves and Symplectic Topology* (2nd ed.) [McS2].

2. Major Topic: Algebraic Topology (Geometry/Topology)

- **Basic Theory:** classical spaces and operations; CW complexes; cellular approximation theorem (sketch); fibrations and cofibrations; homotopy and homotopy equivalence. ([H] Ch. 0)
- **The Fundamental Group:** Seifert–Van Kampen; covering spaces, fundamental group of CW complexes. ([H] Ch. 1)
- **(Co)homology:** singular, simplicial, de Rham, cellular, and Morse (co)homology, Mayer–Vietoris, excision and the LES of a pair; universal coefficient theorem, Künneth formula; multiplications; Poincaré duality; intersection theory. ([H] Ch. 2–3)
- **Higher Homotopy Groups:** LES of a pair and of a fibration; Freudenthal suspension theorem (statement); Whitehead theorem; CW approximation of spaces; Hurewicz theorem; Eilenberg–MacLane spaces; $H^*(X;G) \cong [X, K(G,*)]$; obstruction theory. ([H] Ch. 4)
- **Spectral Sequences:** Leray–Serre and Atiyah–Hirzebruch spectral sequences; multiplicative structure; edge morphisms and transgressions; Gysin sequence; computation of $H^*(SU(n))$, $H^*(\Omega S^n)$, and $\pi_4(S^3)$. ([FF] §21–24, §39)
- **Characteristic Classes:** vector and principal G-bundles, classifying spaces; cohomology of classifying spaces and characteristic classes; Thom isomorphism theorem; characteristic classes as obstructions; Wu formula; splitting principle and basic properties; bordism theory and the Pontryagin–Thom construction. ([MS] §4–15, §17–18)

References. Fomenko and Fuchs, *Homotopical Topology* (2nd ed.) [FF], Hatcher, *Algebraic Topology* [H], Milnor and Stasheff, *Characteristic Classes* [MS].

3. Minor Topic: Partial Differential Equations (Analysis)

- **Some Functional Analysis:** Riesz representation and Lax–Milgram theorems; compact operators; Fredholm alternative; spectral theory of compact self-adjoint operators; Fredholm operators and index. ([Br] Ch. 5–6)
- **Sobolev Spaces:** weak derivatives; extension operators; smooth approximation; traces; Sobolev inequalities and (compact) embeddings. ([Br] §9.1–9.4)
- **Second Order Elliptic PDEs:** existence and uniqueness; elliptic regularity; maximum principle. ([Br] §9.5–9.8 and [Ev] Ch. 6)
- **The Calculus of Variations:** Euler–Lagrange equations; existence of minimizers; regularity; constraints; Lagrange multipliers. ([Ev] Ch. 8)

References. Brezis, *Functional Analysis, Sobolev Spaces and Partial Differential Equations* [Br], Evans, *Partial Differential Equations* (2 ed.) [Ev].

QUAL TRANSCRIPT

GABRIEL BEINER

Committee Members: Ian Agol [A], Michael Hutchings [H], Sung-Jin Oh [O], Fraydoun Rezakhanlou [R].

As people slowly arrived, we discussed the cold weather. I tested out a large collection of whiteboard markers and was dismayed to find none of the purple ones worked. Professor Oh chaired the exam and asked me what topic I would like to begin with. I said algebraic topology.

A: Classify the two-fold coverings of the figure eight $S^1 \vee S^1$.

Me: $\pi_1(S^1 \vee S^1) = \mathbb{Z} * \mathbb{Z}$. There are four two-fold coverings, three of which are connected. However two of them are isomorphic by switching the labelling of the generators of π_1.

A: How did you compute that π_1? What is the definition of isomorphic covers? Are you sure these two are isomorphic?

Me: Seifert–Van Kampen; I gave the argument and Prof. Agol pressed me to be precise on the details. I gave the definition of isomorphic covers. After some prodding I realized the two coverings in question were not isomorphic, since the map interchanging the generators does not commute with projection to the base. The coverings are merely conjugate.

A: How do you know these are all the two-fold covers?

Me: The joining point of $S^1 \vee S^1$ has two preimages in a two-fold cover. By the usual lifting lemma, the two generators in π_1 must lift to either loops or paths between the preimages and the four covers describe all the ways to do that.

A: That is convincing. But how can you see this from the π_1 groups?

Me: I tried to write down generators for the fundamental groups of the covers in terms of the two generators of $\pi_1(S^1 \vee S^1)$. I realized what I initially wrote was wrong and the connected covers require three generators. These presentations lead me nowhere.

The fundamental groups of the two-fold covers should include into $\pi_1(S^1 \vee S^1)$ as index two subgroups whose quotient is the group of deck transformations. But I'm not sure how to classify index two subgroups of $\mathbb{Z} * \mathbb{Z}$.

A: Do such subgroups have nice properties?

Me: They are normal, so we get a short exact sequence $G \to \mathbb{Z} * \mathbb{Z} \to \mathbb{Z}_2$.

A: Okay so such G are classified by maps $\mathbb{Z} * \mathbb{Z} \to \mathbb{Z}_2$ and what are those?

Date: 15 November 2024.

Me: I'm not sure. I don't know group theory.

A: Can you figure out what those maps should be for the covers you have drawn?

Me: The obvious thing is the map should send a generator of $\pi_1(S^1 \vee S^1)$ to either $0, 1$ depending if it lifts to a loop or a path between preimages in the covering. Aha! $\mathbb{Z}_2$ is a very simple group and maps to $\mathbb{Z}_2$ are just determined by if the generators of $\mathbb{Z} * \mathbb{Z}$ map to 0 or 1, and this is what we have just described. The committee seemed to be amused by my delayed elementary observation.

After the exam I figured out how this description works in general. If I have a normal covering of X, the deck transformation group D acts transitively and each transformation is determined by the movement of a single point. Thus there is a surjection $\pi_1(X) \to D$ determined by sending a loop to the deck transformation inducing the corresponding monodromy action. The kernel of this map is exactly π_1 of the cover.

A: Yes, $\mathbb{Z}_2$ is indeed a simple group! Let's move on, is there anything you want to talk about?

Me: We can do whatever. I like homotopy groups.

A: Okay, what is an Eilenberg–MacLane space?

Me: I gave the definition of a $K(G, n)$ and said that they are unique up to WHE.

A: What is the relation between different $K(G, n)$?

Me: The π-group fibration LES gives $\Omega K(G, n) = K(G, n-1)$. Also you can take products.

A: On your syllabus you state $[X, K(G, n)] \cong H^n(X; G)$. Can you explain that for me?

Me: I stated the theorem, explaining how to define a tautological class $F_G \in H^n(K(G, n); G)$ and how pulling it back induces the bijection of cohomology and homotopy classes. I got through most of the proof of surjectivity using a standard constructive argument with obstruction theory. They stopped me before I showed injectivity.

A: You also have spectral sequences on your syllabus, which I don't know so much about. I see you wrote the Gysin sequence, can you explain that?

Me: Sure. (Me, A., and H. discussed how to pronounce Gysin). I wrote down the long exact sequence in homology for a general oriented sphere bundle $\pi : E \to B$ and described what each map is geometrically in the case of a manifold (cup with Euler class, pull-back π^*, and wrong-way map $\pi^!$).

H: You said that the map $H^i(B) \to H^{i+n+1}(B)$ is given by cup product with the Euler class. What does this mean?

Me: I sketched the basics of obstruction theory for sections of fibre bundles and how first obstructions give rise to characteristic classes. For the case of a S^n bundle, this first obstruction lives in $H^{n+1}(B; \mathbb{Z})$. If this bundle is the unit sphere bundle of a rank $n+1$ vector bundle E, then this recovers the usual Euler class $e(E)$.

H: What does this have to do with spectral sequences?

Me: We can consider the Serre spectral sequence of our sphere bundle. I drew the E_2 page.

H: Where does this come from?

Me: I wrote the general expression for the E_2 page of a Serre SS.

H: Should we be using a local coefficient system?

Me: Not if we assume something like homological triviality, which in this case is equivalent to our sphere bundle being oriented. Now note nothing happens in our spectral sequence until the E_n page.

A: That's my page: the "Ian page"! (Actually I should have written the E_{n+1} page which would've spoiled the joke).

Me: The differential here is precisely the cup product with the Euler class as described above, although the proof is a little technical so I won't explain it beyond saying transgressions $H^0(B) \to H^{n+1}(B)$ can usually be described geometrically. Now this is the last page with non-trivial differentials. The two-term short exact sequences we get here patch together to give the Gysin LES. I begin to write the details but H. stops me and says he believes I know what I'm doing and suggests we move on.

A: I was thinking about asking about the Hopf fibration, but I think we can finish here.

The algebraic topology lasted about an hour. We then took a short break. I think everyone took one of the muffins I made, which I was happy about. Professor Oh asked what subject I would like to do next and I said PDE.

R: Saving the best for last!

O: Can you define the Sobolev spaces $W^{k,p}(\mathbb{R}^n)$?

Me: I explain weak derivatives and define the Sobolev spaces as Banach spaces.

O: Now consider a bounded domain $\Omega \subset \mathbb{R}^n$, say with smooth boundary. What is $W^{k,p}(\Omega)$ and how does it relate to the space for $\mathbb{R}^n$?

Me: The definition is basically the same, just restricted to Ω. One can define a bounded extension operator E from Sobolev functions on Ω to ones on $\mathbb{R}^n$. I initially claimed for φ smooth on $\mathbb{R}^n$ this should send $\varphi|_\Omega$ to φ, but realized this is nonsense. I sketched the construction of E in terms of a partition of unity and reflection across the boundary.

O: What do you need to do for $k > 1$? What should you worry about?

Me: I know it can be done but not sure the specifics. I guess you cannot kill things off too quickly for fear of the norms of higher derivatives blowing up.

O: Let's come back to this in a minute. If we are in $W^{k,p}$ for high enough k then we can say something else about our functions. What spaces might these functions live in.

Me: You're talking about Sobolev embedding. The only way I've been able to remember the quantitative parts of this theory is by drawing a graph plotting $1/p$ against k (you can find this image on the Wikipedia page for Sobolev embedding). I then drew this picture. You get embeddings by moving down a line of slope n on this chart. I used this to figure out which p^* gives an embedding $W^{k,p} \hookrightarrow L^{p^*}$ by looking where the line intersects the x-axis.

R: What if the line intersects the y-axis first?

Me: Now you get an embedding into a Hölder space $C^{\ell,\alpha}$. I explained what ℓ and α are in terms of where the intersection is.

O: What are these Hölder spaces?

Me: I wrote the definition of the Hölder spaces. I mistakenly assumed you need to bound the Hölder norm of u and its derivatives, but Prof. Oh pointed out that the latter implies the former. Indeed, if your derivative is bounded then you are automatically Lipschitz.

R: What about the critical case $k = n/p$?

Me: You get an embedding into all L^p spaces for finite p, but not L^∞.

R: Do you get an embedding into any other space?

Me: I guessed bounded mean oscillation (BMO), which I didn't know the meaning of but had seen written in this context.

R: That's right. I'm not sure if this is in Evans.

O: Okay, let's go back to the extension problem. So we see now for large k that the functions have high regularity. What can we say?

Me: I agree, we need our extension to match derivatives. I'm not quite sure how to do that.

O: Let's just simplify to the case of a half-space. (*I still don't know what to do*). Or even simpler a half-line.

Me: Okay given u on $\mathbb{R}^{\leq 0}$, we could extend by reflection $u(-x)$ and then multiply by a bump function φ.

O: Great. So what should we do?

Me: I don't know, maybe choose φ right to match derivatives?

O: How else could we get something that matches with u at zero?

Me: I have no idea. Some kind of polynomial interpolation or Fourier series?

O: There are actually many things you could do. I was thinking of rescaling the reflection $u(-\lambda x)$ and taking a linear combination of these to match derivatives. It's okay you didn't know this, since it's not in Evans' book. Let's move on to elliptic PDE. What are the spaces H_0^1 and H^{-1}?

Me: H_0^1 is the space of $W^{1,2}$ functions which vanish on the boundary in a trace sense. H^{-1} is the dual of H^1, which can be identified with expressions $f - \sum_i \partial_{x_i} f_i$ for $f, f_1, \ldots, f_n \in L^2$. the pairing is given by integration, and using integration by parts to move the derivatives onto the function in H^1.

O+R: It should be the dual of H_0^1.

Me: Sorry, you're right. For the integration by parts to work we need the boundary term to vanish which requires we pair with something in H_0^1.

O: Yes, exactly. Can you tell me what a second order linear elliptic PDE is?

Me: I wrote down a general second order linear differential operator L and gave the uniform ellipticity condition for its second order part.

O: What is a general Dirichlet problem for this operator?

Me: We want to solve $Lu = f$ for $f \in H^{-1}$ with $u \in H_0^1$.

O: Can you tell me about the solvability of this problem? Let's assume throughout all our coefficients are nice and smooth.

Me: Yes, so the weak formulation of this PDE is defined in terms of a bilinear form $B[u, v]$. Should I write this full thing down?

O: Let's say the lower order terms are zero.

Me: I wrote down the corresponding $B[u, v]$. We are looking for u that satisfies $B[u, v] = \langle f, v\rangle$ for all $v \in H_0^1$.

O: How does this relate to a solution of the original problem?

Me: Any classical solution will satisfy this weak solution condition just by integrating by parts. Going back to solvability, the weak solution formula is a Riesz representation-type condition that we can obtain using Lax–Milgram.

O: Can you explain that.

Me: I gave the statement of Lax–Milgram, writing down the continuity and coercivity conditions. Prof. Oh asked me to be specific about the Hilbert spaces I was using. I showed the bilinear form B was continuous and coercive. Here it was critical to invoke the Poincaré inequality for $u \in H_0^1$. Thus we get a unique weak solution for any $f \in H^{-1}$ to our Dirichlet problem.

O: What if L now has non-trivial first order term $\sum b_i(x)\partial_{x_i} u$?

Me: Now B is merely nearly coercive, i.e. it satisfies Gårding's inequality $B[u, v] + \mu\|u\|_{L^2}^2 \geq C\|u\|_{H^1}^2$. We need to apply the Fredholm alternative theory, which says our PDE will either have unique weak solutions for every f or L will have a non-trivial kernel, in which case we get a solution if and only if f is orthogonal to the kernel of the dual problem L^*.

I then explained how we can apply Lax–Milgram to the operator $L_\mu = L + \mu I$. Then Rellich–Kondrachov implies $K = \mu L_\mu^{-1} : H^{-1} \to L^2$ is a compact operator. I explained how solving our original problem is equivalent to finding the kernel/cokernel of $I - K$, which is well understood by the Fredholm alternative. I got a bit confused with the basic algebra relating L and K, and Prof. Oh helped me work it out.

O: Actually in this case where L has no zeroth order term, we are always in the first case of the Fredholm alternative where the Dirichlet problem for $Lu = f$ has a unique weak solution. Let's figure out why. We need to show there are no solutions to $Lu = 0$. What can you say about the properties of a solution u to this problem? I'm looking for a certain phrase.

Me: A solution u must be smooth because of elliptic regularity.

O: Great, let's not go into that. What can you say about the behaviour of a smooth solution of the Dirichlet problem?

Me: We have a maximum principle. The solution must obtain its max and min on the boundary. Since it solves the Dirichlet problem, it must be constantly zero.

O: What are you using about L here?

Me: We need that the zeroth order term is zero. There are weaker statements if the zeroth order term is non-zero but with a fixed sign.

O: Okay, that's all for PDEs.

This section again lasted a little less than an hour. We took another short break and began symplectic.

H: What possible genus can an embedded symplectic submanifold Σ in $\mathbb{C}P^2$ have?

Me: Hmm... I don't know about this. Are we assuming Σ is oriented?

H: Well it needs to be symplectic.

Me: Of course, whoops. Um, maybe we can use the fact the normal bundle of Σ is identified with its cotangent bundle. What am I saying... that's for Lagrangian not symplectic submanifolds!

H: Everyone always asks about Lagrangian submanifolds of $\mathbb{C}P^2$, so I decided to ask about symplectic ones!

Me: Okay, the Fubini–Study form is Poincaré dual to the homology class of a hyperplane. Maybe we should think about what homology class can be represented by Σ?

H: That seems like a good idea.

Me: Uh, I'm still not sure. I know that $\mathbb{C}P^1 \subset \mathbb{C}P^2$ is an example for genus zero, maybe this is the only possibility.

H: Can you get any other examples, maybe from another field of math?

Me: Aha, we could take a degree d algebraic curve $\mathbb{C}P^1 \to \mathbb{C}P^2$ and by the degree-genus formula it would have genus $g = (d-1)(d-2)/2$.

H: Where does this come from? Is this true for any symplectic surface?

Me: I think so, we can use adjunction.

> *I didn't recall until after the exam that every symplectic surface is a J-holomorphic curve for some J and so this is clearly true by the adjunction formula in that context. In fact this question should have all been trivial to me, since I had given a talk a couple weeks earlier on minimal genus problems where I stated the symplectic Thom conjecture, which solves this problem immediately.*

H: Can you explain?

Me: Okay let's say the fundamental class $[\Sigma]$ in homology equals $d[\mathbb{C}P^1]$.

H: What can d be here?

Me: I think it must be positive. We have $\int_\Sigma \omega_{FS} = d\int_{\mathbb{C}P^1} \omega_{FS}$ and d must be positive since Σ needs to have positive area.

I then wrote down the standard short exact sequence of vector bundles for an embedding and took first Chern classes to obtain $c_1(T\mathbb{C}P^2|_\Sigma) = c_1(\nu_\Sigma) + c_1(T\Sigma)$. I could not figure out what the left hand side should be.

H: We're doing algebraic topology again! What is the right hand side?

Me: We have $c_1(\nu_\Sigma) = [\Sigma]^2 = d^2[\mathbb{C}P^1]^2 = d^2$. This follows because $\mathbb{C}P^1 \subset \mathbb{C}P^2$ has normal bundle the Hopf bundle, with self-intersection number one. Also we know that $c_1(T\Sigma) = \chi(\Sigma) = 2 - 2g$.

H: Okay, what should the left hand side $c_1(T\mathbb{C}P^2|_\Sigma)$ be when $d = 1$?

Me: If $d = 1$ then $g = 1$ and the left side should be 3 for the formula to work (I struggled with this simple algebra for a minute). So I guess it should be $3d$ in general.

I still didn't understand why $3d$ *was right but no one asked me to explain. The next day I realized that* $c_1(T\mathbb{C}P^2) = 3[\mathbb{C}P^1]$ *from the usual formula for Chern classes of projective space, and by naturality, this is multiplied by* d *under pullback by inclusion, i.e. restriction.*

H: Okay great, so symplectic surfaces must have genus of the form $(d-1)(d-2)/2$. Did you need that the surfaces were symplectic, or is this true for any embedded submanifold?

Me: I know it cannot be true in general for a silly reason: otherwise the Thom conjecture would make no sense. But I'm not sure where we used symplectic-ness. Maybe in the computation of $c_1(T\mathbb{C}P^2|_\Sigma)$?

H: Well, what did you need about your bundles to take the first Chern classes of the exact sequence?

Me: Hmm, they should be oriented... Oh, this needs to be an exact sequence of complex vector bundles! And this only makes sense if Σ is something like a complex or symplectic submanifold.

H: Great, let's move on. On your syllabus you have the phrase contact type. What is a contact type hypersurface?

Me: It's a hypersurface in a symplectic manifold (M, ω) with a contact form which is a primitive of ω. Equivalently, it is a hypersurface with a locally defined transverse Liouville vector field.

H: Can you prove these two definitions are equivalent?

Me: If I have a transverse Liouville vector field X, then $\iota_X\omega$ defines a one-form α. I began the computation that α was contact on the hypersurface, contracting with X to show $\omega^n \neq 0$ implies $\alpha \wedge \mathrm{d}\alpha^{n-1} \neq 0$, but Hutchings stopped me half of the way through.

H: I think we get the idea. What about the other direction?

Me: I know the idea is to identify a neighbourhood of the hypersurface with the symplectization.

H: What is a symplectization?

Me: Given (Y, α) contact, the symplectization is $(\mathbb{R}_s \times Y, \mathrm{d}(e^s\alpha))$, which is symplectic.

Back to the problem, I should be able to find a zero-section preserving diffeomorphism from a tubular neighbourhood of the hypersurface Y to $\mathbb{R} \times Y$, but this requires me to know Y has trivial normal bundle. I'm not sure why that's true. Let me skip that for a moment. Given the diffeomorphism, this clearly preserves the symplectic form along the hypersurface Y. Hence by the Moser isotopy lemma, we can find a diffeomorphism between these neighbourhoods preserving Y and pulling back one symplectic form to another.

H: Okay, let's fill in the gap. What do you need for the normal bundle to Y to be trivial?

Me: It's just a line bundle, so we need an orientation. Ah, but Y and M are both oriented from the contact/symplectic forms and so this normal bundle will be too.

H: And finishing the argument, what is the Liouville vector field?

Me: You just pushforward ∂_s from the symplectization.

H: Anyone else have anything to ask?

R: What is Gromov non-squeezing?

Me: I stated the theorem.

R: Can you prove it? How many proofs do you know?

Me: I can do it with pseudoholomorphic curve theory. I think there is also a proof with symplectic homology.

I was thinking of the proof of the Weinstein conjecture for convex domains. But I think there may be a proof with the Viterbo symplectic homology capacity.

R: Also (Hofer–Zehnder) capacities. Okay, let's see Gromov's original proof.

Me: I sketched the outline of the proof, starting from compactifying the cylinder to $S^2 \times T^{2n-2}$ and finding a tame almost-complex structure extending the one pushed forward from the embedded ball. I explained that one wants to find a J-holomorphic sphere passing through the centre of the ball in the homology class $[S^2 \times \mathrm{pt}]$.

To show such a thing existed, I wrote down the corresponding moduli space of simple genus zero J-curves with one marked point. I proved what the moduli space was for the standard product complex structure (diffeomorphic to $S^2 \times T^{2n-2}$ via the evaluation map) and said that this choice of complex structure is regular by automatic transversality. I said that if J is also regular, a homotopy will give a cobordism of moduli spaces and thus a J-curve whose marked point hits the centre of the embedded ball, since the cobordism preserves the degree of the evaluation map. If J is non-regular, we can find a sequence of regular J's limiting to it and the corresponding sequence of curves converges to the curve we want. These last two facts required Gromov compactness of our moduli space.

R: What is Gromov compactness? How do you know we have compactness in this case?

Me: I stated that $W^{1,p}$ bounds for $p > 2$ gives C^∞ convergence by bootstrapping, but energy bounds only gives $W^{1,2}$ bounds. This is the critical number in Sobolev embedding and so that sequence may diverge in L^∞. But, the possible singularities are very mild and correspond after reparametrization to sphere bubbles at points where the L^∞ norm blows up. I explained how the topological energy considerations force there to be finitely many bubbles. Thus the space of J-spheres is compactified by the introduction of trees of genus zero nodal curves. In our case, $[S^2 \times \mathrm{pt}]$ is a primitive cohomology class, and the energy of such a curve cannot split into two bubbles each of smaller but positive energy. Hence the moduli spaces in question are already compact.

Now we can finish the proof. Given the desired J-curve, its pullback under the ball embedding gives a minimal surface in the ball. Its area is bounded below by the area of 2-disk in the ball using the monotonicity formula. Noting the homology class of the J-sphere, the symplectic area of S^2 in the compactified cylinder is also bounded below by the symplectic area of the 2-disk. Writing this explicitly gives the desired inequality in Gromov non-squeezing.

No one had any more questions. The symplectic section was surprisingly short, maybe 40 minutes. The whole exam, with breaks, was about 2 and a half hours. They asked me to step outside for about 30 seconds after which they congratulated me on passing and everyone shook my hand. Hutchings told me to enjoy a relaxing weekend.

Overall, the exam went quite smoothly and was even somewhat enjoyable. My nerves went away early on and I felt most of the questions were on things I had been well prepared for. The committee was quite friendly and I never felt that worried about embarrassing myself.

I was surprised that the symplectic section was so short, I think shorter than PDEs. I knew to expect it, but I was still surprised at how many syllabus topics were not discussed/tested at all. I was not asked to compute any (co)homology, π groups, or anything about classifying vector bundles. I was not asked anything about the proofs of Sobolev inequalities, elliptic regularity, etc., or to make arguments with challenging inequalities or Green's identities computations, and calculus of variations was not mentioned. I was not asked anything about Hamiltonian actions and moment maps, examples of symplectic manifolds, Hamiltonian dynamics, Moser tricks, or to detail any analysis around J-holomorphic curves. Still, I am happy the threat of having to discuss all of these topics loomed over me, as it made me learn them much better in preparation for the exam!