LARGE CARDINALS, FORCING AND REFLECTION

JAMES CUMMINGS

These notes reflect, fairly faithfully, the contents of a minicourse which I delivered in September 2013 at the RIMS symposium on "Reflection principles and the set theory of large cardinals". They retain the informal style and rapid pace of the original four lectures, in particular many proofs are only sketched or are done from much stronger hypotheses than are really needed.

I would like to record here my warmest thanks to Professor Sakaé Fuchino, the organiser of the symposium, for his kind invitation and for his excellent hospitality during my first visit to Japan. Professor Fuchino was also kind enough to send me his detailed record of my lectures, which was invaluable in the preparation of these notes.

Further reading: The reader who would like to learn more about the material treated in these lectures is referred to my surveys of singular cardinal combinatorics [2] and forcing and large cardinals [3], Eisworth's survey of singular cardinal combinatorics [4], Foreman's comprehensive survey of the theory of generic elementary embeddings [6], and Kanamori's book on large cardinals [8].

1. LECTURE ONE

1.1. **Reflection and non-reflection.** A very useful way of bringing some order to the diverse models, theories, principles and axioms appearing in modern combinatorial set theory is to use the notion of "reflection" as an organising principle. To make this point clearer, I will start by classifying some well-known ideas under the broad headings of "Reflection" and "Non-reflection".

Reflection	Non-reflection
Compactness	Incompactness
Stationary reflection	Non-reflecting stationary sets
Large cardinal axioms	$V = L$
The tree property	Aronszajn trees
Chang's conjecture	Jensen's $\square$ principle

Mediating between these various reflection and non-reflection principles is the theory ZFC. "ZFC combinatorics" is a rich area including for example Shelah's PCF theory and Todorčević's theory of minimal walks. We note that ZFC proves results of both reflection type (eg Silver's theorem on cardinal arithmetic, Shelah's singular compactness theorem) and non-reflection type (eg the existence of an $\aleph_1$-Aronszajn tree, Shelah's results on partial squares and the existence of stationary sets in $I[\lambda]$, strong negative partition properties of $\aleph_1$).

1.2. **Large cardinals and elementary embeddings.** We recall the definitions of some standard large cardinals in terms of elementary embeddings.

- κ is measurable if and only if there is an elementary embedding $j : V \to M$ with $\mathrm{crit}(j) = \kappa$. In this case there exists such an embedding with $^\kappa M \subseteq M$.
- κ is λ-supercompact if and only if there is an elementary embedding $j : V \to M$ with $\mathrm{crit}(j) = \kappa$, $j(\kappa) > \lambda$ and $^\lambda M \subseteq M$.

- κ is supercompact if and only if it is λ-supercompact for every λ.
- κ is huge if and only if if there is an elementary embedding $j : V \to M$ with $\mathrm{crit}(j) = \kappa$ and $^{j(\kappa)}M \subseteq M$.

1.3. From large cardinals to reflection.

Reflection principles often have high consistency strength, and so it is not surprising that they are closely connected with large cardinals. In particular, large cardinal axioms have many interesting consequences of reflection type. Typically large cardinal axioms will only give reflection properties for large structures, and the major theme in subsequent lectures will be the question of how to "miniaturise" and get reflection properties for small structures.

For the sake of simplicity, many of the arguments we give use large cardinal hypotheses stronger than necessary.

A) Stationary reflection:

    i) Stationary reflection for ordinals: If $\kappa = \mathrm{cf}(\kappa) > \omega$ and $S \subseteq \kappa$ is stationary, then S *reflects* if there is $\alpha < \kappa$ such that $\mathrm{cf}(\alpha) > \omega$ and $S \cap \alpha$ is stationary in α.

    If κ is measurable then every stationary subset of κ reflects. To see this fix j with critical point κ, and observe that $j(S) \cap \kappa = S$. Since $M \subseteq V$ we see that S is stationary in M, so $M \models \exists \alpha < j(\kappa) \, j(S) \cap \alpha$ is stationary. We are done by elementarity.

    We note that the hypothesis that κ is weak compact would suffice. If $V = L$ then for every regular uncountable κ which is not weakly compact, there is a stationary set $S \subseteq \kappa$ which does not reflect [7].

    ii) A highly significant generalisation: let X be uncountable and let $S \subseteq [X]^{\aleph_0}$. S is *stationary* iff S meets every subset of $[X]^{\aleph_0}$ which is club (cofinal and closed under unions of increasing ω-sequences), or equivalently for every structure $\mathcal{M}$ for some countable first order language L there is $A \in S$ with $A \prec \mathcal{M}$.

    A stationary set $S \subseteq [X]^{\aleph_0}$ *reflects* to an uncountable set $X_0 \subseteq X$ if $S \cap [X_0]^{\aleph_0}$ is stationary in $[X_0]^{\aleph_0}$.

    Let κ be λ-supercompact and let $S \subseteq [\lambda]^{\aleph_0}$ be stationary. Fix $j : V \to M$ witnessing the λ-supercompactness of κ, then we claim that in M the stationarity of $j(S)$ reflects to $j``\lambda$; by elementarity, the stationarity of S reflects to some $X_0 \subseteq \lambda$ such that $|X_0| < \kappa$ and $X_0 \cap \kappa \in \kappa$.

    To see the claim, fix $\mathcal{M} \in M$ some structure with underlying set $j``\lambda$. Copying in the obvious way via $j$, we obtain a structure $\mathcal{M}_0$ on $\lambda$, and so there is $A \in S$ with $A \prec \mathcal{M}_0$. Now $j(A) = j``A \in j(S)$ and $j(A) \prec \mathcal{M}$.

B) The tree property: Let T be a κ-tree with underlying set κ, and let j have critical point κ. Then $j(T) \restriction \kappa = T$, and fixing some point on level κ in $j(T)$ and looking at points below we see that T has a branch in M. Since $M \subseteq V$, T has a branch in V.

    We just showed that κ being measurable implies that κ has the tree property. This can be viewed as a reflection argument, in which we showed that "T has no cofinal branch" reflects.

    As in part A i) above, the hypothesis that κ is weak compact would suffice. Also as in part A i), if $V = L$ then κ has the tree property if and only if κ is weakly compact [7].

C) Chang's conjecture: A structure $\mathcal{M}$ for a countable first order language with a distinguished predicate U is said to be a (λ, κ)-*structure* if the underlying set of

$\mathcal{M}$ has size λ and $U^{\mathcal{M}}$ is a subset of size κ. The general form of the Chang Conjecture, written $(\lambda_1, \kappa_1) \twoheadrightarrow (\lambda_0, \kappa_0)$ asserts "every (λ_1, κ_1)-structure has a (λ_0, κ_0) elementary substructure".

Let κ be huge with a witness $j : V \longrightarrow M$ such that $j(\kappa) = \lambda$. Let $\mathcal{M}$ be a (λ, κ)-structure with underlying set λ and $U^{\mathcal{M}} = \kappa$. Then in M we have that $j(\mathcal{M})$ is a $(j(\lambda), \lambda)$-structure and $j``\lambda \prec j(\mathcal{M})$, what is more $j``\lambda$ has order type $\lambda = j(\kappa)$ and $j``\lambda \cap j(\kappa) = \kappa < j(\kappa)$. Reflecting we get that there is a $(\kappa, < \kappa)$ substructure, that is (in the obvious notation) $(\lambda, \kappa) \twoheadrightarrow (\kappa, < \kappa)$.

D) Compactness for the chromatic number: Let κ be measurable and G a graph with underlying set κ, let $\mu < \kappa$ and let $\chi(G) > \mu$ (that is, there is no μ-colouring of the vertices in which adjacent vertices receive distinct colours). We claim there is an induced subgraph H with $|H| < \kappa$ and $\chi(H) > \mu$.

The proof is easy; if $j : V \longrightarrow M$ has critical point κ and G has underlying set κ then the subgraph which $j(G)$ induces on κ is just G, and since $M \subseteq V$ it is clear that in M we have $\chi(G) > \mu$.

E) Compactness for non-existence of transversals: Let κ be measurable and let $\mathcal{H}$ be a family of countable sets where $|\mathcal{H}| = \kappa$ and $\mathcal{H}$ has no transversal (a transversal is a 1-1 choice function). Then there is $H' \subseteq H$ such that $|H'| < \kappa$ and H' has no transversal. The proof is left to the reader.

F) PCF theoretic example: Let λ be singular with $\mathrm{cf}(\lambda) = \mu$. By a basic fact in PCF, there exist a sequence $\langle \lambda_i : i < \mu \rangle$ of regular cardinals which is increasing and cofinal in κ and a sequence $\langle f_\zeta : \zeta < \lambda^+ \rangle$ which is increasing and cofinal in $\prod_{i < \mu} \lambda_i$ under the eventual domination ordering $<^*$.

If $\eta < \lambda^+$ is limit then an *eub (exact upper bound)* for $\langle f_\zeta : \zeta < \eta \rangle$ is a function h such that $f_\zeta <^* h$ for all $\zeta < \eta$, and $g <^* h$ implies there is $\zeta < \eta$ with $g <^* f_\zeta$. An eub may not exist, but if it does it is easily seen to be unique modulo eventual equality.

Now suppose there is κ such that $\mu < \kappa < \lambda_0 < \lambda$, and κ is λ^+-supercompact. Consider the ordinal $\rho = \sup j``\lambda^+$ and the function $H$ given by $H(i) = \sup j``\lambda_i$. It is routine to check that in M, H is an eub for the image of the scale under j. Reflecting, we see that there is a point $\eta < \lambda^+$ with $\mu < \mathrm{cf}(\eta) < \kappa$ such that there is an eub h for $\langle f_\zeta : \zeta < \eta \rangle$ where $\mathrm{cf}(h(i))$ is greater than μ and increases strictly with i.

With more work we can show that there is a singular cardinal $\nu < \kappa$ such that stationarily many η of cofinality ν^+ have an eub as above. This is a very strong failure of square: it contradicts the weak square principle $\square_\mu^*$, and also the even weaker property "μ^+ is approachable".

G) (A severely weakened form of) the Rado conjecture, whose full version is discussed later.

Say that a tree T of height ω_1 is *special* if there is a function f from T to ω such that $s <_T t \implies f(s) \neq f(t)$, so that in particular T has no cofinal branch. Suppose that κ is measurable and T is a non-special tree of height ω_1 with underlying set κ. Then the usual reflection arguments show that T has a non-special subtree of size less than κ.

2. Lecture Two

In the last lecture we saw several examples in which κ is a large cardinal, and an elementary embedding with critical point κ is used to get reflection properties for structures

of size at least κ. In this lecture we turn to the question of getting similar reflection properties for small structures. It is natural to consider elementary embeddings with small critical points in this connection, but we recall from Lecture One that the critical point of an embedding $j : V \longrightarrow M \subseteq V$ must be a measurable cardinal.

The key idea for getting reflection for small structures is to use *generic elementary embeddings*. These are elementary embeddings of the general form $j : V \longrightarrow M \subseteq V[G]$ where j and M are defined in $V[G]$. The good news is that such embeddings can have small critical points such as $\aleph_1^V$. The bad news is that typically $M \not\subseteq V$, and this will mean that most of the reflection arguments from the end of the first lecture do not easily generalise. We will discuss two (closely related) ways of obtaining a generic elementary embedding, namely lifting and generic ultrapowers.

2.1. Lifting.

Let M and N be two transitive models of ZFC and let $\pi : M \longrightarrow N$ be elementary. Let $\mathbb{P} \in M$ be a forcing poset. If G is $\mathbb{P}$-generic over M and H is $\pi(\mathbb{P})$-generic over N, then we may hope to lift π to obtain an elementary embedding $\pi^+ : M[G] \longrightarrow N[H]$.

Silver gave a sufficient criterion for π^+ to exist, namely that $\pi``G \subseteq H$. Assuming this, we attempt to define π^+ by $\pi^+ : i_G^M(\dot{\tau}) \mapsto i_H^N(\pi(\dot{\tau}))$. The definition makes sense by an easy argument appealing to the Definability and Truth lemmas plus the elementarity of π: if $i_G^M(\dot{\tau}) = i_G^M(\dot{\sigma})$ then there is $p \in G$ such that $p \Vdash_{\mathbb{P}}^M \dot{\sigma} = \dot{\tau}$, and then $\pi(p) \in H$ and $\pi(p) \Vdash_{\pi(\mathbb{P})}^N \pi(\sigma) = \pi(\dot{\tau})$, so that finally $i_H^N(\pi(\dot{\tau})) = i_H^N(\pi(\dot{\sigma}))$. Similar arguments show that π^+ is elementary, and it is easy to see (using canonical names for ground model sets) that π^+ extends π.

We note that if π^+ is defined in this way then $\pi^+(G) = H$. Conversely if there exists π^+ which extends π and maps G to H, then $\pi``G \subseteq H$.

Silver's criterion raises an obvious question: how to ensure that the condition $\pi``G \subseteq H$ is satisfied? We will say that $q$ is a *strong master condition* (for $\pi$ and $G$) if $q \leq \pi(p)$ for all $p \in G$. If this is so then we may force to get $H$ which is $\mathbb{Q}$-generic over $N$ with $H \ni q$, and any such $H$ will have the property that $\pi``G \subseteq H$.

2.2. An easy lifting argument.

We give what is probably the simplest interesting example of a lifting argument. No master condition will be required here.

Let κ be measurable and let $j : V \longrightarrow M$ be an embedding with critical point κ and $^\kappa M \subseteq M$. Let $\delta = \mathrm{cf}(\delta) < \kappa$, and let $\mathbb{P} = Coll(\delta, < \kappa)$. Let g be $\mathbb{P}$-generic over V.

We pause to analyse $j(\mathbb{P})$, where by elementarity we have $j(\mathbb{P}) = Coll(\delta, < j(\kappa))^M$. Since $^\kappa M \subseteq M$, this is just $Coll(\delta, < j(\kappa))$. $j(\mathbb{P})$ is isomorphic in the obvious way to $\mathbb{P} \times \mathbb{Q}$, where $\mathbb{Q}$ is the poset of all partial functions $q : \delta \times [\kappa, j(\kappa)) \longrightarrow j(\kappa)$ such that $|q| < \delta$ and $q(\eta, \zeta) < \zeta$ for all $(\eta, \zeta) \in \mathrm{dom}(q)$. In a mild abuse of notation we write $\mathbb{Q} = Coll(\delta, [\kappa, j(\kappa)))$, and identify $j(\mathbb{P})$ with $\mathbb{P} \times \mathbb{Q}$.

Each $p \in \mathbb{P}$ is a partial function of size less than δ, so that easily $j(p) = (p, 1_\mathbb{Q})$ If we now force over $V[g]$ with $\mathbb{Q}$ we obtain a filter h such that $g \times h$ is $j(\mathbb{P})$-generic over V (hence over M) and $j``g \subseteq g \times h$. So working in $V[g * h]$ we may lift j and obtain $j^+ : V[g] \longrightarrow M[g * h]$. This is an example of a generic embedding of $V[g]$ obtained by forcing over $V[g]$ with $\mathbb{Q}$. To lighten the notation we will also denote the extended embedding by j. It is important to note that forcing with g added no $< \delta$-sequences of ordinals, so that $\mathbb{Q} = Coll(\delta, [\kappa, j(\kappa)))^{V[g]}$.

We now give an example to illustrate that not all the reflection arguments from Lecture One go through for generic embeddings. Our example involves the tree property. Let $\delta = \omega$, so that $\kappa = \aleph_1^{V[g]}$, and let $T \in V[g]$ be an Aronszajn tree. As before $j(T) \upharpoonright \omega_1 = T$,

and so there is a cofinal branch $b \in M[g * h]$. But since $M[g * h] \not\subseteq V[g]$ we are not warranted in concluding that $b \in V[g]$, and indeed no such branch can exist in $V[g]$.

On the other hand some reflection arguments can be rescued, at least in a limited form. Suppose now that $\delta = \omega_1$, so that $\kappa = \omega_2^{V[g]}$. We claim that in $V[g]$, every stationary subset of $\omega_2 \cap \mathrm{cof}(\omega)$ reflects (this is as much reflection as we can hope for, since no stationary subset of $\omega_2 \cap \mathrm{cof}(\omega_1)$ can reflect).

Let $S \in V[g]$ be a stationary subset of $\omega_2 \cap \mathrm{cof}(\omega)$. As usual $j(S) \cap \kappa = S$, and we note that in $M[g * h]$ we have $\mathrm{cf}(\kappa) = \omega_1$. We will be done if we can show that S is stationary in $M[g * h]$.

To see this, start by observing that S is named by a canonical $\mathbb{P}$- name in M since $V_{\kappa+1} \subseteq M$, and hence $S \in M[g]$. Since $M[g] \subseteq V[g]$ and S is stationary in $V[g]$, it is also stationary in $M[g]$. Now h is generic over $M[g]$ for the countably closed forcing $\mathbb{Q}$ and so by a standard proper forcing argument S is stationary in $M[g * h]$.

Generalisation: let κ be supercompact and as in the previous example let $\mathbb{P} = Coll(\omega_1, < \kappa)$ with g which is $\mathbb{P}$-generic. Let $\lambda \geq \kappa$ be regular and let $S \subseteq \lambda \cap \mathrm{cof}(\omega)$ be stationary .

Let j witness the λ-supercompactness of κ, and argue as above that S is stationary in $M[g * h]$. This implies (with some work, the key point is that j is continuous at points of small cofinality) that $j(S) \cap \sup(j``\lambda)$ is stationary in $M[g * h]$. It follows by the usual argument that in $V[g]$ every stationary set reflects to a point of cofinality ω_1.

Essentially the same argument gives a stronger reflection principle in the same model: For every stationary set $S \subseteq [\lambda]^{\aleph_0}$ there is a set $X \subseteq \lambda$ such that $|X| = \aleph_1 \subseteq X$ and S reflects to X. We return to this topic in Lecture Four.

2.3. Generic ultrapowers.

Let $X \in V$ and suppose that by forcing we can introduce an ultrafilter U on the V-powerset of X. In $V[U]$ we may form an ultrapower of V using functions in V with domain X, we call this $Ult(V, U)$. In general $Ult(V, U)$ will not be well-founded, but when it is we may collapse it and obtain a model M and a generic embedding $j : V \to M \subseteq V[U]$.

The typical way to obtain a suitable U is to force with the poset consisting of I-positive sets for some ideal I on X. This gives an ultrafilter U such that $U \cap I = \emptyset$. The ideal I is said to be *precipitous* if it is forced that $Ult(V, U)$ is well-founded.

There is a close connection between the theory of lifting and the theory of precipitous ideals. This topic is explored in great detail in Foreman's survey paper on generic embeddings [6], and we content ourselves here with a couple of remarks related to the easy lifting example that we gave above. Recall that κ is measurable, $\delta < \kappa$ is regular, and g is generic for $Coll(\delta, < \kappa)$.

- There is a precipitous ideal on δ^+ in $V[g]$, namely $I = \{X \subseteq \kappa : \Vdash_{\mathbb{Q}}^{V[g]} \kappa \notin j(X)\}$.
- In $V[g]$, $P(\kappa)/I$ has a dense δ-closed subset.

3. LECTURE THREE

Recall from the last lecture the Silver criterion for lifting an elementary embedding $\pi : M \longrightarrow N$ between transitive models of set theory to a map $\pi^+ : M[G] \longrightarrow N[H]$, where G is generic over M for $\mathbb{P} \in M$ and H is generic over N for $\pi(\mathbb{P})$. If $\pi``G \subseteq H$ then this lifting is possible. We record a few easy remarks about this situation, and the related concept of "strong master condition".

(1) If $\pi``G \subseteq H$ then in fact $G = \pi^{-1}``H$, so that H determines G. To see this suppose for contradiction that we have $p \in (\pi^{-1}``H) \setminus G$. By the truth lemma there is $p' \in G$ forcing that $\hat{p} \notin \dot{G}$, that is to say that p' is incompatible with p.

By elementarity $\pi(p)$ is incompatible with $\pi(p')$, which is absurd since both $\pi(p)$ and $\pi(p')$ lie in H.

(2) Suppose that q is a strong master condition for π and G, that is $q \leq \pi"G$. Then
   (a) $G = \pi^{-1}"H$ for any N-generic H such that $H \ni q$.
   (b) For any $A \in M$ which is a maximal antichain of $\mathbb{P}$, there is a unique $q' \in \pi"A$ such that $q \leq q'$.

3.1. More on master conditions.

We will say that q is a strong master condition for π if condition 2b) above holds. Suppose that q has this property and that $H \ni q$ is an N-generic filter. It is easy to check that if $G = \pi^{-1}"H$ then G is an M-generic filter, and also $G = \{p : q \leq \pi(p)\}$.

In summary, q is a strong master condition if it forces that the preimage of the N-generic filter is an M-generic filter, and it also determines the preimage. We say that q is a weak master condition if it forces the preimage of the N-generic to be M-generic: one can show that this is equivalent to a weakening of 2b) above, stating that for any $A \in M$ which is a maximal antichain of $\mathbb{P}$ the set $\pi"A$ is predense below q. The heart of the matter is that a weak master condition forces that $\pi"A$ meets the $\mathbb{Q}$-generic filter while a strong master condition decides exactly where $\pi"A$ meets it.

3.2. Master conditions and proper forcing.

Let $X \prec H_\theta$ be countable, let M be the collapse of X and let $\pi : M \longrightarrow H_\theta$ be the inverse of the collapsing map. Of course in general H_θ and M are not models of the full ZFC axioms, but they are models of enough set theory to make the lifting arguments and the discussion of master conditions as above go through for any poset $\mathbb{P} \in M$.

If we let $\mathbb{Q}$ be a forcing poset in X and $\mathbb{P}$ be its collapse, so that $\pi(\mathbb{P}) = \mathbb{Q}$, then the definitions of strong and master weak conditions are familiar concepts from the theory of proper forcing.

To be more explicit:

- q is a weak master condition for π iff for every maximal antichain A of $\mathbb{P}$ with $A \in M$ the set $\pi"\mathbb{A}$ is predense below q iff for every maximal antichain B of $\mathbb{Q}$ with $B \in X$ the set $B \cap X$ is predense below q iff q is $(X, \mathbb{Q})$-generic.
- Similarly, q is a strong master condition for π iff for every maximal antichain B of $\mathbb{Q}$ with $B \in X$ there is $b \in B \cap X$ with $q \leq b$ iff q is totally $(X, \mathbb{Q})$-generic.

3.3. Master conditions in action: Laver indestructibility.

We will sketch a notable theorem of Laver [9]: if κ is supercompact then by set forcing we can make the supercompactness of κ indestructible under κ-directed closed forcing.

We need another theorem of Laver (which we will not prove) asserting the existence of a *Laver diamond*. If κ is supercompact then there exists a function $f : \kappa \to V_\kappa$ such that for all sets $x \in V$ and all cardinals $\lambda \geq \kappa$ there is $j : V \longrightarrow M$ such that j witnesses the λ-supercompactness of κ and $j(f)(\kappa) = x$.

Now for the indestructibility. We define an iteration of length κ with Easton supports using a function f as above. We will define iterands $\mathbb{Q}_\alpha$ together with ordinals $\lambda_\alpha < \kappa$ as follows: $\mathbb{Q}_\alpha$ is trivial forcing and $\lambda_\alpha = 0$ unless $f(\alpha) = (\dot{\mathbb{Q}}, \lambda)$, where $\dot{\mathbb{Q}}$ is a $\mathbb{P}_\alpha$-name for a κ-directed closed forcing poset, λ is an ordinal and $\lambda_\eta < \lambda$ for all $\eta < \alpha$. In this case we set $\lambda_\alpha = \alpha$ and $\dot{\mathbb{Q}}_\alpha = \dot{\mathbb{Q}}$.

Let $\mathbb{P} = \mathbb{P}_\kappa$, where standard arguments show that $\mathbb{P}$ is κ-cc forcing of size κ. Let G be $\mathbb{P}$-generic and let $\mathbb{Q} \in V[G]$ be a κ-directed closed forcing poset, say $\mathbb{Q} = i_G(\dot{\mathbb{Q}})$. Let $\mu_0 \geq \kappa$, with a view to showing that κ is μ_0-supercompact in $V[G]$.

Fix μ_1 which is much larger than μ_0. Exactly how large it needs to be can be found by inspecting the proof. We can afford to be careless here since we have the strong hypothesis that κ is supercompact. Now use the guessing property of the Laver diamond function f to find j witnessing that κ is μ_1-supercompact with $j(f)(\kappa) = (\dot{\mathbb{Q}}, \mu_1)$.

In M, $j(\mathbb{P}_\kappa)$ is an iteration computed using the function $j(f)$. It is easy to see that $j(f) \restriction \kappa = f$, and that V and M agree on the iteration of length κ computed from f. So $j(\mathbb{P}_\kappa)$ can be factored as $\mathbb{P} * \mathbb{Q} * \mathbb{R}$, where $\mathbb{R}$ is the part of the iteration between κ and $j(\kappa)$.

Let g be $\mathbb{Q}$-generic over $V[G]$. We chose $\mu_1 > |\mathbb{P} * \mathbb{Q}|$, so an easy argument involving names shows that $V[G * g] \models {}^{\mu_1}M[G * g] \subseteq M[G * g]$. By the rules governing the iteration and the fact that $j(f)(\kappa) = (\dot{\mathbb{Q}}, \mu_1)$, all iterands in $\mathbb{R}$ are at least μ_1^+-closed and so by standard facts about Easton iterations $\mathbb{R}$ is μ_1^+-closed in $M[G * g]$, hence it is μ_1^+-closed in $V[G * g]$.

Now we force over $V[G * g]$ with $\mathbb{R}$ and obtain a generic object H. Since every $p \in G$ has bounded support in κ, $j``G \subseteq G * g * H$ and we may lift to obtain $j : V[G] \longrightarrow M[G * g * H]$. We may argue that $j``g \in M[G * g * H]$ (using the fact that this model can see g and very large initial segments of the lifted j).

Now $|j``g| < j(\kappa)$ and this is a directed subset of the forcing $j(\mathbb{Q})$, which by elementarity is $j(\kappa)$-directed closed in $M[G * g * H]$. So we may find $q \in j(\mathbb{Q})$ with $q \leq j``g$, and force over $V[G * g * H]$ with $j(\mathbb{Q})$ to obtain $h \ni q$. Since $q \leq j``g$ we have $j``g \subseteq h$, so we may lift again to get $j : V[G * g] \longrightarrow M[G * g * H * h]$.

It is important to understand that we are not quite done! The lifted j is a generic elementary embedding which can only exist in a generic extension of $V[G * g]$. To finish, we define an ultrafilter U on $(P_\kappa \mu_0)^{V[G*g]}$ by $A \in U \iff j``\mu_0 \in j(A)$. Since μ_1 is much bigger than μ_0, we may use the closure of $\mathbb{R} * j(\mathbb{Q})$ and argue that $U \in V[G * g]$, where it will serve as a witness to the μ_0-supercompactness of κ.

Note: If κ is inaccessible then a κ-*Kurepa tree* is a tree of height κ such that level α has size at most $|\alpha| + \aleph_0$. It is easy to see that if κ is measurable then no such tree can exist. There is a κ-closed forcing poset to add such a tree, so that it is not possible to make supercompactness indestructible under κ-directed closed forcing.

3.4. Saturated ideals.

Another (closely related) way of obtaining generic embeddings with strong reflection properties involves saturated ideals. Let μ be regular and uncountable, then a normal ideal I on μ is *saturated* iff $P(\mu)/I$ has the μ^+-cc.

Solovay showed that if we force with $P(\mu)/I$, we obtain an ultrafilter U on the V-powerset of μ such that $Ult(V, U)$ is wellfounded: what is more, if we form the transitive collapse of $Ult(V, U)$ we get a model M and a generic embedding $j : V \to M \subseteq V[U]$ with the additional property that $V[U] \models {}^\mu M \subseteq M$.

In the special case where $\mu = \delta^+$ we see that cardinals up to δ are preserved in $V[U]$, and that $j(\mu) = (\mu^+)^V = (\delta^+)^V$, so that adding U collapses μ but preserves μ^+.

Kunen gave a forcing argument that the cardinal $\aleph_1$ can carry a saturated ideal. The starting hypothesis is that there is a huge cardinal κ. Say we have $j : V \longrightarrow M$ with $\mathrm{crit}(j) = \kappa$, $j(\kappa) = \lambda$, ${}^\lambda M \subseteq M$. In the course of the construction κ will become $\aleph_1$ and λ will become $\aleph_2$.

Here is a very brief sketch of Kunen's construction. We force with $\mathbb{P} * \dot{\mathbb{Q}}$ where $\mathbb{P}$ is κ-cc forcing of size κ which collapses κ to become the new ω_1, and $\mathbb{Q}$ is countably closed forcing of size λ which collapses λ to become the new ω_2. The iteration is designed so

that $j(\mathbb{P})$ absorbs $\mathbb{P} * \mathbb{Q}$, and the closure of M implies that the quotient forcing is λ-cc in $V^{\mathbb{P}*\dot{\mathbb{Q}}}$.

Let G be $\mathbb{P}$-generic over V and let H be $\mathbb{Q}$-generic over $V[G]$. Forcing over $V[G*H]$ with $j(\mathbb{P})/G*H$ we prolong $G*H$ to a generic object G^*, and we lift to get $j : V[G] \longrightarrow M[G^*]$. The forcing $\mathbb{Q}$ is chosen so that the union of $j``G_{\mathbb{Q}}$ is a condition in $j(\mathbb{Q})$, so we may use this union as a master condition and obtain $j : V[G * H] \longrightarrow M[G^* * H^*]$. Finally we may use the countable closure of $j(\mathbb{Q})$ and the λ-cc of $j(\mathbb{P})/G * H$ to read off a suitable ideal in $V[G * H]$.

4. Lecture Four

In this lecture we first discuss some relations between large cardinals, forcing axioms and reflection principles. We conclude with a discussion of the role of "preservation theorems" in arguments involving generic embeddings, and give a few examples.

4.1. The weak reflection principle.
As we saw in an earlier lecture, if κ is supercompact and we force with $Coll(\aleph_1, < \kappa)$ then we obtain a model of the following "Weak reflection principle" introduced by Foreman, Magidor and Shelah [5].

Note: We use the term "weak reflection principle", which is now common in the literature but is not the terminology of [5], to distinguish WRP from a stronger principle which appears later in the lecture. The WRP is in fact a very strong principle, whose consequences are explored in [5].

WRP: For all $\lambda \geq \aleph_2$ and all stationary $S \subseteq [\lambda]^{\aleph_0}$, there is $X \subseteq \lambda$ such that $\aleph_1 = |X| \subseteq X$, and $S \cap [X]^{\aleph_0}$ is stationary

If $\pi : \aleph_1 \to X$ is a bijection, then for any $A \subseteq [X]^{\aleph_0}$ it is easy to see that A is stationary iff $\{\alpha : \pi``\alpha \in A\}$ is stationary in ω_1. What is more this stationary subset is (modulo clubs) independent of the choice of π. So $S \cap [X]^{\aleph_0}$ in the statement of WRP is "morally" a stationary subset of ω_1.

One notable consequence of WRP proved in [5]: If WRP holds then every *stationary-preserving forcing* (that is every forcing which preserves stationary subsets of ω_1) is semiproper.

4.2. The consistency of MM.
Martin's Maximum (MM) is the axiom which states that for every stationary preserving $\mathbb{P}$ and every family of $\aleph_1$ many dense sets in $\mathbb{P}$, there is a filter meeting each dense set in the family. The statement SPFA is obtained by replacing "stationary preserving" by "semiproper". Foreman, Magidor and Shelah [5] showed that MM is consistent relative to a supercompact cardinal, and we outline their proof here. The main idea is that the natural model of SPFA is a model of WRP, hence a model of MM.

Let κ be supercompact and let f be a "Laver diamond" function as above. We use f to guess names for semi-proper forcing posets in a revised countable support (RCS) iteration $\mathbb{P}$ of length κ. Since there are many stages where we add Cohen reals and many other stages where we collapse the local ω_2, it is not hard to see that $\kappa = \omega_2 = 2^\omega$ in $V[G]$. Let G be $\mathbb{P}$-generic.

Claim 1: SPFA holds in $V[G]$.

Proof: Let $\mathbb{Q}$ be a semiproper poset in $V[G]$ and let $\vec{D}$ be a list of ω_1 dense sets in $V[G]$. Let μ be very large, let j be an embedding witnessing that κ is μ-supercompact with $j(f)(\kappa) = \dot{\mathbb{Q}}$ some name for $\mathbb{Q}$. One can verify that $\mathbb{P}$ is an initial segment of $j(\mathbb{P})$. Let h be $\mathbb{Q}$-generic over V and let H be $j(\mathbb{P})/G * h$-generic over $V[G * h]$. By properties

of RCS iterations, $\mathbb{P}$ is defined as a direct limit and so as above we have $j"G \subseteq G * h * H$ and may form a lifted embedding $j : V[G] \longrightarrow M[G * h * H]$.

By choosing μ large enough we may arrange that $j"h \in M[G * h * H]$. As h is generic it meets each dense set listed along $\vec{D}$. Since $\text{crit}(j) = \omega_2$ and $\vec{D}$ has length ω_1, $j"h$ meets each set listed along $j(\vec{D})$. So easily $j"h$ generates a filter which meets each dense set along $j(\vec{D})$, and we are done by a familiar appeal to elementarity.

Claim 2: WRP holds in $V[G]$.

Proof: Let $S \in V[G]$, let $S \subseteq [\lambda]^{\aleph_0}$ be a stationary set in $V[G]$ and let μ be much larger than λ. Choose j witnessing that κ is μ-supercompact with $j(f)(\kappa)$ naming the poset $Coll(\aleph_1, \lambda)$. Let $h : \aleph_1 \to \lambda$ be generic for this collapse poset, and let $T = \{\alpha : h"\alpha \in S\}$. As in Claim 1 we may find a suitable H and lift j to obtain $j : V[G] \longrightarrow M[G * h * H]$.

S is stationary in $V[G * h]$ because the collapse poset is proper, so T is stationary in $V[G * h]$ and hence in $M[G * h]$. By properties of RCS iterations, H is generic for semi-proper forcing and so T is stationary in $M[G * h * H]$. Hence S is stationary in $M[G * g * H]$, and we finish as in previous reflection arguments.

4.3. The strong reflection principle.

Another result from [5] shows that MM implies WRP. It is interesting to note that PFA, although a very strong axiom, does not imply any amount of stationary reflection: Beaudoin [1] and independently Magidor showed that we can force over any model of PFA by a forcing which preserves PFA and adds a non-reflecting stationary subset of $\omega_2 \cap \text{cof}(\omega)$. PFA does imply a rather subtle form of reflection, namely Moore's MRP (Map Reflection Principle) [11].

After the initial work on MM, Shelah [12] showed that SPFA implies MM. Reflecting on Shelah's proof, Todorčević extracted a stronger reflection principle SRP (Strong Reflection Principle) which follows from SPFA.

SRP: Let $S \subseteq [\lambda]^{\aleph_0}$ (not necessarily stationary). Let θ be a regular cardinal much larger than λ and let $X \in H_\theta$. There is a continuous increasing $\in$-chain $\langle M_i : i < \omega_1 \rangle$ of countable elementary submodels of H_θ such that $X \in M_0$, and for all i we have that $M_i \cap \lambda \in S$ iff there is $N \prec H_\theta$ with $M_i \cap \omega_1 = N \cap \omega_1$, $M_i \subseteq N$ and $N \cap \lambda \in S$.

SRP implies WRP, and also that NS is saturated. The main trick for using it: given a chain $\vec{M}$ as above, find $N \prec H_\theta$ countable with $\vec{M} \in N$, let $i = N \cap \omega_1$ and observe that $M_i \subseteq N$ with $M_i \cap \omega_1 = N \cap \omega_1$.

To illustrate the trick in action, we sketch the proof that SRP implies WRP. Let $S \subseteq [\lambda]^{\aleph_0}$, let θ be much larger than λ, and let $\langle M_i : i < \omega_1 \rangle$ be a chain of models as in the conclusion of SRP with $S \in M_0$. Now let $X = \bigcup_i M_i \cap \lambda$: we claim that $S \cap [X]^{\aleph_0}$ is stationary, which amounts by a previous discussion to claiming that $T = \{i : M_i \cap \lambda \in S\}$ is stationary in ω_1. If not let C be disjoint from T, and build N such that $C, \vec{M} \in N$ and $N \cap \lambda \in S$. Let $i = N \cap \omega_1$ and observe that $i \in C$, however we have that $N \cap \lambda \in S$ so that $M_i \cap \lambda \in S$, hence $i \in T$.

4.4. Preservation theorems.

Suppose that we wish to derive some reflection principle in V from the existence of a generic elementary embedding $j : V \longrightarrow M \subseteq V[G]$. The typical situation will be something like this: we have an structure X in V with some property P, and would like to assert that there is a small substructure X_0 with P. To do this we will consider $j"X$ (typically this will be a member of M and a substructure of $j(X)$) and attempt to argue that $j"X$ has P in M. Often this is done by arguing that X still has P in $V[G]$, using special properties of the forcing $\mathbb{P}$ for which G is generic.

In the arguments that we gave so far P has typically been some form of stationarity: the main preservation theorems that we used were two results by Shelah, "proper forcing preserves stationary subsets of $[\lambda]^{\aleph_0}$ for all uncountable λ" and "semi-proper forcing preserves stationary subsets of $\aleph_1$". We conclude these lectures by discussing some other preservation theorems and associated reflection results.

4.4.1. The Rado conjecture.
The Rado Conjecture asserts that every non-special tree of height ω_1 has a non-special subtree of cardinality $\aleph_1$. Todorčević proved that if κ is supercompact and we force with $Coll(\omega_1, < \kappa)$ then in the extension $V[G]$ the Rado conjecture holds. Here is a sketch of the argument.

Let $T \in V[G]$ be a tree of height ω_1 such that every subtree of cardinality $\aleph_1$ is special. For a suitable generic embedding $j : V[G] \longrightarrow M[G * H] \subseteq V[G * H]$ we may take it that $j``T \in M[G * H]$ and $|j``T| = \aleph_1$, where H is generic over $V[G]$ for the countably closed forcing poset $Coll(\omega_1, < j(\kappa))$. It follows readily that T is special in $V[G * H]$.

Preservation theorem: If T is special in a countably closed forcing extension then T is special.

Proof: Let $\mathbb{P}$ be countably closed and let $\dot{g}$ name a specialising function. Build a family $\langle p_t, \alpha_t : t \in T \rangle$ where p_t forces that $g(t) = \alpha_t$, and $s <_T t \implies p_t \leq p_s$. It is easy to see that $t \mapsto \alpha_t$ is a specialising function.

4.4.2. The tree property.
Starting with a weakly compact cardinal κ, Mitchell showed that by doing a suitable κ-cc forcing $\mathbb{M}$ we can obtain a model in which $2^\omega = \kappa = \omega_2$ and κ still has the tree property.

Here is a very rough sketch of the main idea, working for simplicity on the assumption that κ is measurable. Let G be $\mathbb{M}$-generic and let $\mathbb{Q} = j(\mathbb{M})/G$. If we force H which is $\mathbb{Q}$-generic over V, then we may lift to obtain $j : V[G] \longrightarrow M[G * H]$ and argue as in Lecture 1 that $T \in M[G]$ and T has a branch b in $M[G * H]$. Every proper initial segment of b has the form $\{s : s <_T t\}$ for some $t \in T$, in particular every initial segment of b is in $M[G]$. The key point is that the poset $\mathbb{Q}$ satisfies the following preservation theorem in $M[G]$:

Preservation theorem: If $x \subseteq \kappa$ is in the $\mathbb{Q}$-generic extension and every proper initial segment of x is in the ground model, then x is in the ground model.

4.4.3. Stationary reflection.
Magidor [10] showed it to be consistent that every stationary subset of $\aleph_{\omega+1}$ reflects. His proof produces a model V in which for every n with $0 < n < \omega$ there is a generic embedding $j : V \longrightarrow M \subseteq V[G]$ such that $\text{crit}(j) = \omega_n$, $j``\aleph_{\omega+1} \in M$, and G is generic for ω_{n-1}-closed forcing. It is tempting to think this is enough to get the desired result, but there is a snag. We would hope to use the following "fact":

Preservation theorem (FALSE IN GENERAL): If $\lambda < \mu$ are regular cardinals and $\mathbb{P}$ is λ^+-closed then $\mathbb{P}$ preserves all stationary subsets of $\lambda \cap \text{cof}(\mu)$.

Shelah showed that if GCH holds and κ is supercompact, then there is a generic extension in which for some stationary set $S \subseteq \aleph_{\omega+1} \cap \text{cof}(\omega_1)$, the stationarity of S can be destroyed by $\aleph_2$-closed forcing. The argument, which we omit involves a "miniaturisation" of the PCF reflection fact from Lecture One.

The argument of [10] requires an extra step in which the model is shown to have an extra combinatorial property ("$\aleph_{\omega+1}$ has the approachability property"). We can then use a theorem of Shelah:

Preservation theorem: If $\lambda < \mu$ are regular cardinals, μ has the approachability property, and $\mathbb{P}$ is λ^+-closed then $\mathbb{P}$ preserves all stationary subsets of $\lambda \cap \mathrm{cof}(\mu)$.

REFERENCES

[1] R. Beaudoin, The proper forcing axiom and stationary set reflection, Pacific Journal of Mathematics 149 (1991), 13-24.

[2] J.Cummings, Notes on singular cardinal combinatorics (survey), Notre Dame Journal of Formal Logic 46 (2005), 251–282.

[3] J. Cummings, Iterated forcing and elementary embeddings. Handbook of Set Theory, M. Foreman and A.Kanamori eds, 775–883, Springer, 2010.

[4] T. Eisworth, Successors of Singular Cardinals. Handbook of Set Theory, M. Foreman and A.Kanamori eds, 1229–1350, Springer, 2010.

[5] M. Foreman, M. Magidor and S. Shelah, Martin's maximum, saturated ideals, and nonregular ultrafilters I, Annals of Mathematics 127 (1988), 1-47.

[6] M. Foreman, Ideals and generic elementary embeddings. Handbook of Set Theory, M. Foreman and A.Kanamori eds, 885–1147, Springer, 2010.

[7] R. Jensen, The fine structure of the constructible hierarchy, Annals of Mathematical Logic 4 (1972), 229-308

[8] A. Kanamori, The higher infinite. Springer, 1994.

[9] R. Laver, Making the supercompactness of κ indestructible under κ-directed closed forcing, Israel Journal of Mathematics 29 (1978), 385–388.

[10] M. Magidor, Reflecting stationary sets, Journal of Symbolic Logic 47 (1982), 755–771.

[11] J. T. Moore, Set mapping reflection, Journal of Mathematical Logic 5 (2005), 87-98.

[12] S. Shelah, Semiproper Forcing Axiom implies Martin Maximum but not PFA$^+$, Journal of Symbolic Logic 52 (1987), 360–367.