Very Short Course on Forcing
for Participants of
XIV International Congress of Topology

I'd like to show you that the basic idea of forcing is really quite natural and simple. In addition, the apprentice forcer needn't be, in my view, concerned with the more basic technical aspects (Cohen's theorem). I'd like to convince you that you could be reading forcing arguments in a very short time.

The major obstacle to forcing is that we let our Platonic view of the set-theoretic universe cause us to doubt the possibility that forcing works. We must remember that the language of set-theory is simply $\in$ and that we are really studying consequences of the axioms and not the strucure of the "real" universe of sets.

Through some logical formalism we can take the following view: There are countable transitive models, M, of ZFC. If we fix any infinite set (element) P of M, there are sets $G \subset P$ which are not in M. We construct an extension $M[G]$ of M which includes G. $M[G]$ should be as small as possible and be a model of ZFC.

The model $M[G]$ is *generated* by names.

Fix any set $X \in M$ and any set $\tau \subset P \times X$ with $\tau \in M$. τ is a P-*name* for a subset of X.

Since, τ, X and G are all in $M[G]$, it is very clear that we must also have, in $M[G]$,

$$\tau_G = \{x \in X : (\exists p \in G)\ (p, x) \in \tau\}$$

We also say that τ is a name for τ_G. I.e. τ_G simply must be in $M[G]$ and τ is how it gets there.

Forcing is the technique of making use of the fact that all subsets of X in $M[G]$ have such a description from elements of M.

Again, for emphasis, each subset of X in $M[G]$ is of the form τ_G for some $\tau \subset P \times X$ with $\tau \in M$.

But clearly we need more things in $M[G]$, sets of new subsets of X and so on. E.g. we can fix a collection $\{\tau_y : y \in Y\} \in M$ and then $\{(\tau_y)_G : y \in Y\}$ is again in $M[G]$. It's a good technical exercise that each subset of $\mathcal{P}(X)$ has such a name.

The official names are recursive: if Σ is a collection of P-names and $\sigma \subset P \times \Sigma$, then σ is a P-name (and we can only use those that are members of M). But we won't need them.

Next, we don't use just any P and G but rather restrict to the case that P is a poset and then G will be a *generic* subset.

This is how we can gain some control over $M[G]$.

Poset will mean atomless, separative and with a largest element (1). Separative means if $p \not\leq q$, then there is an $r < p$ such that $r \perp q$. $r \perp q$ means no element is below both.

$D \subset P$ is *dense* if each $p \in P$ is above some member of D.

$G \subset P$ is a filter if it is upwards closed and downwards directed (given $p, q \in G$, there is an $r \in G$ below both)

$G \subset P$ is P-generic (over a model M) if G meets every dense subset of P which is a member of M. I always think of P as a $\downarrow$-tree and G as a new branch. Let $\mathcal{D}_P$ denote the set of dense subsets of P. Since M is countable, $\mathcal{D}_P \cap M$ is countable, so there are such G (this is the only place we need the fact that M is countable).

It is a little confusing to realize that all countable atomless posets are forcing equivalent, so it is really the family $\mathcal{D}_P \cap M$ which is the key. But it is best not to think about this; just take the view that M is the whole universe of sets and that by some magic a generic filter G can be found. Actually, it is necessary (and obvious) to realize that each p in P is a member of many generic filters.

Lemma: G will not be P-generic over $M[G]$ (equivalently, there are no members of M which are P-generics over M).

Proof. (Identifying dense subsets of P).

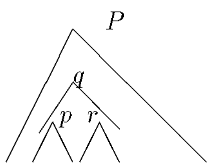

Check that if G is a filter, then
$D_G = \{p \in P : p^{\downarrow} \cap G = \emptyset\}$ is dense.
To see this, let $q \in P$ and show that p or r is in D_G.
Since G is directed, it contains something below at most one of p and r, hence one of them is in D_G.

Here is a name for D_G: $\{(p, q) : p \perp q\}$. It is good to think this through.

Some useful notions related to *dense*:

$A \subset P$ is pre-dense below a condition p, if $A^{\downarrow} = \bigcup\{a^{\downarrow} : a \in A\}$ is dense in $p^{\downarrow}$. In terms of meeting G, A and $A^{\downarrow}$ are indistinguishable.

p is called *stronger* than q if $p < q$ since $p \in G$ implies $q \in G$ (more information about where G will go).

$A \subset P$ is an *antichain* if $a \perp b$ for all distinct $a, b \in A$. Every maximal antichain is pre-dense. For each set A, there is an antichain B, such that $A^{\downarrow} = B^{\downarrow}$.

The cellularity (chain condition) of P is the minimum cardinal such that each antichain is smaller. Of course, ccc posets are quite commonly mentioned. This condition *preserves* ω_1 which is crucial for proving $\neg$CH is consistent.

Nice names are a useful concept and every name τ has an equivalent *nice* name. For a name τ and a value $x \in X$, set

$$A_x = \{p \in P : (p, x) \in \tau\}$$

In general, replace x by a "smaller than τ" name σ. It is *nice* if each A_x is an antichain.

If $(p, x) \in \tau$, then

(*) $(\forall G)$ (generic) $p \in G$ implies $x \in \tau_G$; say that $p \Vdash x \in \tau$.

$p \Vdash x \in \tau$ (i.e. (*)) iff A_x is pre-dense below p.

This is because A_x pre-dense below p iff $p \Vdash G \cap A_x \neq \emptyset$.

On the other hand, we'd say that $p \Vdash x \notin G$ if there's no generic G such that $p \in G$ and $x \in \tau_G$. This simply means that $A_x^{\downarrow} \cap p^{\downarrow}$ is disjoint. This is different from $p \not\Vdash x \in \tau$. We say that p decides the statement $x \in \tau$ if either $p \Vdash x \in \tau$ or $p \Vdash x \notin \tau$. The set $\{p \in P : p \text{ decides } x \in \tau\}$ is dense.

This is not discussing any particular G and can result in important information, in M, about the behaviour of τ in the "generic" generic extension. Consider the case, for example, when P is countable and X is uncountable. Some particular p will make the '$x \in \tau$' decision for uncountably many $x \in X$.

I often use the notation τ_p and τ_p^+ for

$$\{x : p \Vdash x \in \tau\} \ \text{ and } \ \{x : (\exists q < p) q \Vdash x \in \tau\}$$

These sets are both in M and $p \Vdash \tau_p \subset \tau \subset \tau_p^+$. Neither involves generic filters or forcing, just looking at the various A_x's. E.g., τ_p^+ is just those x such that $p^{\downarrow}$ meets $A_x^{\downarrow}$.

For a formula φ, $p \Vdash \varphi(\tau_1, \ldots, \tau_n)$ means that

$$M[G] \models \varphi((\tau_1)_G, (\tau_2)_G, \ldots, (\tau_n)_G)$$

holds for all generic G such that $p \in G$.

The set of those p which decide $\varphi(\tau_1, \ldots, \tau_n)$ is dense. This innocuous seeming statement has the following very important re-formulation:

Forcing Theorem: For each statement φ and generic filter G, if $M[G] \models \varphi$, then there is a $p \in G$ such that $p \Vdash \varphi$.

Nothing happens in $M[G]$ "by accident", we can figure it out from M by looking at names and dense sets - sometime called "by genericity".

It is important to remember that if $q < p$ and $p \Vdash \varphi$ then $q \Vdash \varphi$.

A basic example: force with $2^{<\omega_1}$, then

1. ω_1 is preserved
2. the old $\mathfrak{c}$ now has cardinality ω_1
3. the old $\mathfrak{c}^+$ is now the cardinal ω_2
4. CH holds
5. $\mathbb{R}^M = \mathbb{R}$
6. $\diamondsuit$ holds.
7. there are new subsets of ω_1 but not too many.

Certain things passing from M to $M[G]$ are obvious. non-cardinals do not become cardinals, but some cardinals may not remain cardinals. Basic statements, e.g. $x \in y$ do not change.

In the jargon, this poset is countably closed and has the $\mathfrak{c}^+$-chain condition.

countably closed is guaranteeing that ω_1 is preserved and that no new reals are added.

The chain condition is guaranteeing that $\mathfrak{c}^+$ remains a cardinal. Also the new value of 2^{ω_1} is bounded by $(|P| \cdot \omega_1)^{<c(P)}$ (count the nice names).

Suppose $\tau \subset P \times (\mathfrak{c} \times \mathfrak{c}^+)$. This τ is a potential map witnessing that $\mathfrak{c}^+$ does not remain a cardinal. We may assume that it is a nice name and that 1 forces that it is a function (you may wish to check out the *maximum principle* in Kunen's text).

For each $\gamma < \mathfrak{c}$, and $\xi, \zeta < \mathfrak{c}^+$, $A_{(\gamma,\xi)}$ and $A_{(\gamma,\zeta)}$ (as per above) are antichains of P. But, in this case, the members of $A_{(\gamma,\xi)}$ and $A_{(\gamma,\zeta)}$ are pairwise incompatible because $\tau(\gamma)$ cannot equal both ξ and ζ (which is equivalent to assuming that 1 forced that it was a function).

Since $\bigcup\{A_{(\gamma,\xi)} : \xi \in \mathfrak{c}^+\}$ has cardinality at most $\mathfrak{c}$, there is a set, $F(\gamma)$, with cardinality at most $\mathfrak{c}$ such that $\tau(\gamma)$ is a member of $F(\gamma)$. I.e., $F(\gamma)$ is just the set of those ξ such that $A_{(\gamma,\xi)}$ is not empty. The main point to our discussion till now is to realize that if $\zeta \notin F(\gamma)$, then there is no G such that $M[G] \models \tau(\gamma) = \zeta$.

The function F is a ground model approximation to the function τ (It is essentially τ_1^+). The set $\bigcup_\gamma F(\gamma)$ has an upper bound (computed in M) and thus, so does τ. $(\mathfrak{c}^+)^M$ remains bigger than $|(\mathfrak{c})^M|$.

P is countably closed if each countable chain has a lower bound.

Lemma: a countably closed poset does not add any new countable sets of ordinals.

Proof: Suppose that τ is a P-name of a function from ω into M (and $\tau \in M$). We will show that there is a condition p which decides the value of $\tau(n)$ for each n. Since p and τ are members of M, so is this sequence of values (i.e., τ_p). Thus p forces that τ is equal to a ground model function. I.e. for each n, let $h(n)$ denote the unique value such that $p \Vdash \tau(n) = h(n)$. Thus $p \Vdash \tau = h$. This is just $\tau_p = h = \tau_p^+$.

It is trivial to find p: Start with any p_0. Since the set D_0 of conditions that decide the value of $\tau(0)$ is dense, p_0 has an extension $p_1 \in D_0$, i.e. which decides the value of $\tau(0)$. Similar p_1 has an extension, p_2 in D_1, and continue. So p_n can be chosen for each n (NO assumptions on P used yet). Now use P is countably closed to pick a lower bound p.

We started with any p_0 and showed it had an extension, p, deciding τ is in M, thus the set of such p is dense and so every generic must meet it. I'll be careless with this point, most arguments should start below a general $p \in P$ since $M[G] \models ??$ means there is a p forcing the situation.

CH holds: for each $\alpha \in \omega_1$, we can decode a real from G. Note that $\bigcup G$ is just a function from ω_1 to 2. Thus $\bigcup G \restriction [\alpha, \alpha + \omega)$ can be thought of as a function from ω to 2 (the α'th). In this way, we can think of G as supplying us with a list of ω_1 many reals. We prove every real is listed by simply showing (exercise) that for each $r \in {}^\omega 2$,

$$D_r = \{p \in 2^{<\omega_1} : (\exists \alpha) p(\alpha + n) = r(n) \text{ for each } n\}$$

is dense.

Here's the proof of $\diamondsuit$. For each limit $\alpha < \omega_1$, set $a_\alpha \subset \alpha$ equal to $\{\beta \in \alpha : (\cup G)(\alpha + \beta) = 1\}$.

Assume $p_0 \Vdash \tau \subseteq \omega_1$ and $\dot{C} \subset \omega_1$ is a cub. It suffices to find a $q < p_0$ and an α such that $q \Vdash \alpha \in \dot{C}$ and $\tau \cap \alpha = a_\alpha$.

Construct $p_0 > p_1 > p_2 \cdots$ together with $\alpha_n = \mathrm{dom}(p_n)$ and β_n (n odd) so that, e.g.,

p_2 decides $\tau \cap \alpha_1$, $\alpha_2 < \beta_3$
$p_3 \Vdash \beta_3 \in \dot{C}$ (as for each $2n$ and $2n+1$).

Then $p_\omega = \bigcup_n p_n$ decides $\tau \cap \alpha$ where $\alpha = \sup_n \alpha_n$ and forces that α is in $\dot{C}$ (a little tricky since it uses that $\dot{C}$ is a cub).

Thus, $q = p_\omega \frown \chi_{\tau \cap \alpha}$ is the desired condition.

With these, we've seen some standard basic arguments. Others are similar but more subtle combinatorics on P are required for more delicate distinctions in $M[G]$.

Now we look at 4 posets that add reals

For discussion, let $\mathcal{F}$ denote a set of increasing functions in ω^ω (e.g. a scale) and let $\mathcal{A}$ denote any maximal (descending) tower in $\mathcal{P}(\omega)$. I.e. for all $b \in [\omega]^\omega$, there's an $A \in \mathcal{A}$ with $|A \cap b| = \omega$ and assume for technical reasons that $\mathcal{A}$ is closed under finite intersections, $\mathcal{F}$ is closed under finite sups, and both are closed under finite modifications.

$P_C = {}^{<\omega}\omega$

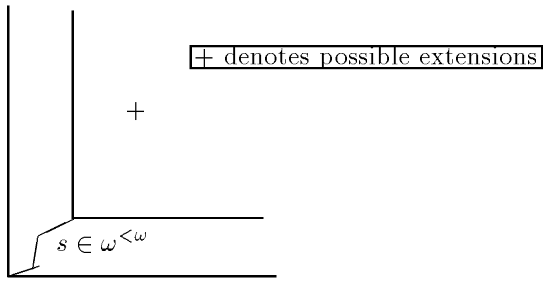

$P_{\mathcal{F}} = \{(s, f) : f \in \mathcal{F} \text{ and } s \subset f\}$
ordered by $(s, f) \leq (t, g)$ if $g \leq f$.

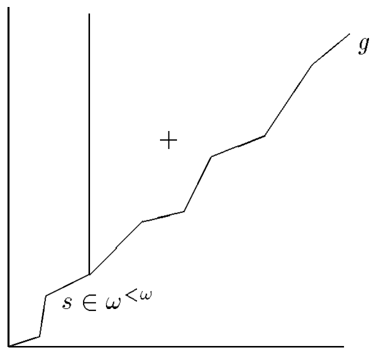

$P_{\mathcal{A}} =$

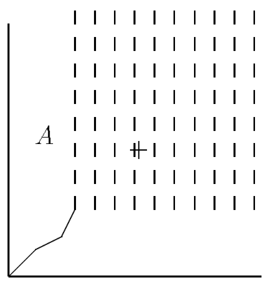

where $P_{\mathcal{A}} = P_C \times \mathcal{A}$
and $(t, B) \leq (s, A)$ if
$B \subset A$, $s \subset t$ and $\operatorname{ran}(t \setminus s) \subset A$

Finally let $P_{\mathcal{M}}$ denote the measure algebra, i.e. compact positive measure subsets of $[0, 1]$ ordered by reverse inclusion.

All are ccc and it is easy to see how each naturally adds a real.

e.g. if G is a generic for $P_{\mathcal{M}}$, then for each n, exactly one of $[i/2^n, (i+1)/2^n]$ will be in G.

None of them have a violent effect on M but there are differences:

we may preserve that $\mathcal{F}$ is dominating; or only preserve unbounded but add an unbounded real; or add a real which dominates $\mathcal{F}$.

Similarly for $\mathcal{A}$: we may or may not add a pseudointersection for $\mathcal{A}$; i.e. maximality of $\mathcal{A}$ may or may not change.

In the following, G will denote a generic filter for the poset under discussion.

Fact. For $P_{\mathcal{F}}$: $h = \bigcup\{s : (\exists f)(s, f) \in G\}$ will be bigger, mod finite, than every member of $\mathcal{F}$.

Proof: $D_f = \{(s, g) \in P_{\mathcal{F}} : f \leq^* g\}$ is dense. Remember our picture? A condition (s, g) forces that $h(n) \geq g(n)$ for all $n \notin \operatorname{dom}(s)$.

Fact. For P_C: $\bigcup G$ witnesses that $\mathcal{F}$ is not dominating in $M[G]$. It can jump real high any time it wants - conditions can't say much.

Again, for each f,

$$D_f = \{s \in {}^{<\omega}\omega : (\exists n) s(n) \geq f(n)\}$$

is dense.

Fact. For P_C: If $\mathcal{F}$ is unbounded in M then $\mathcal{F}$ remains unbounded in $M[G]$. Just as it was hard to get $\bigcup G$ to stay small, it's also hard to get it to stay big.

Proof: Let $\dot{h}$ be a P_C-name of a member of ω^ω. We must produce a condition p and an $f \in \mathcal{F}$ such that $p \Vdash f \not\leq^* \dot{h}$.

Enumerate P_C as $\{p_n : n \in \omega\}$. For each n, fix $g(n)$ such that

$$(\forall k \leq n)(\exists q \leq p_k) q \Vdash \dot{h}(n) \leq g(n)\} \ .$$

Find $f \in \mathcal{F}$ such that $f \not<^* g$. Suppose there is some p_ℓ and some j such that $p_\ell \Vdash \dot{h}(n) \geq f(n)$ for $n \geq j$. It suffices to find $q \leq p_\ell$, $q \Vdash \dot{h}(n) < f(n)$.

Pick any n larger than j and ℓ such that $f(n) > g(n)$. By the definition of $g(n)$, there is a $q \leq p_\ell$ such that $q \Vdash \dot{h}(n) \leq g(n) < f(n)$.

Fact. For $P_{\mathcal{M}}$: In $M[G]$, each $h \in \omega^\omega$ is bounded by some $f \in M$. (refine infinitely often)

Let $\dot{h}$ be a name in ω^ω. For each integer n,
$D_n = \{K \in P_{\mathcal{M}} : (\exists m)\ K \Vdash \dot{h}(n) = m\}$ is dense. (i.e. each $(K, (n, m))$ is a member of $\dot{h}$)

For each n, pick a finite subset $\mathcal{K}_n$ of D_n such that $\mu\left(\bigcup \mathcal{K}_n\right)$ is greater than $1 - \frac{1}{2^{n+2}}$. Then $K = \bigcap_n \left(\bigcup \mathcal{K}_n\right)$ has measure $\geq 1/2$

For each n, there is a value m, call it $f(n)$, such that $\bigcup \mathcal{K}_n \Vdash \dot{h}(n) \leq f(n)$.
Indeed, the finite antichain $\mathcal{K}_n$ is predense below $\bigcup \mathcal{K}_n$ and, a priori, K. I.e., K forces that G picks a member of $\mathcal{K}_n$.

Thus $K \Vdash \dot{h}(n) \leq f(n)$ for all n.

Finally we look at the tower $\mathcal{A}$:

Easy: $P_{\mathcal{A}}$ adds b such that $b \setminus a$ is finite for all $a \in \mathcal{A}$, namely $\bigcup G''\omega$; from our picture; $(s, A) \in G$ implies $\bigcup G \setminus s[\omega]$ is contained in A.

Medium: P_C does not add such a b

Hard: $P_{\mathcal{F}}$ also does not add such a b (Probably needs that $\mathcal{F}$ is dominating)

But for more interesting consistency results (beyond $V \neq L$) we must add more than one real. To do this effectively (not adding the reals, but prove the results) we must discuss iterated forcing.

Iterated forcing is really a recursive process for building posets. E.g. $P = \bigcup_\alpha P_\alpha$. After some λ steps, we've got P_λ. Well, we can investigate the universe obtained by forcing with P_λ and then decide that more forcing should be done. So we cleverly build a $P_{\lambda+1} \supset P_\lambda$ (and continue).

At some stage $\kappa > \lambda$, we stop; and we learn that $V[G_\kappa]$ can be regarded as a forcing extension of $V[G_\lambda]$.

That is, there is a poset, denoted P_κ/P_λ in the model $V[G_\lambda]$, such that there is a generic filter $H \subset P_\kappa/P_\lambda$ such that

$$V[G_\kappa] = V[G_\lambda][H]$$

May say $H = G_\kappa/G_\lambda$ - factoring the forcing.

Even when we can construct a simple (but large) poset P all at once it is still very valuable to be able to view P as an iteration of various P_0 followed by P/P_0.

The correct notation is $P * Q$ for forcing with P followed by forcing with Q.

But note, we are just building a poset P_κ, and factoring the model $V[G_\kappa] = V[G_\lambda][G_\kappa/G_\lambda]$.

Valuable? $V[G]$ is hard to investigate, but if we are interested in an object X, there may be $P_0 \subset P$ such that X already lives in $V[G \cap P_0]$. We can determine what properties X has in $V[G\cap P_0]$ and then investigate which properties are preserved by the forcing P/P_0.

(Models of set theory are better understood than names of names).

The down side is that to learn what is preserved by P/P_0 we still have to get into the nitty-gritty of the iterated construction of posets.

Let us first look at the single stage viewed as an iteration. The poset is $\mathrm{Fn}(\omega_2,2)$ (Cohen reals): $\mathrm{Fn}(I,2)$ is simply finite partial functions from I into 2 ordered by reverse inclusion. We will assume CH holds in V.

It happens that $\mathrm{Fn}(I\dot\cup J,2)=\mathrm{Fn}(I,2)*\mathrm{Fn}(J,2)$

Theorem of $V[G]$: If $\mathcal{F}\subset\omega^\omega$ is unbounded then it has an $\aleph_1$-sized subset which is also unbounded.

The plan: find $P_0\subset P$, $P_0=\omega_1$, set $\mathcal{F}'$ equal $\mathcal{F}\cap V[G\cap P_0]$ so that $\mathcal{F}'$ is unbounded in $V[G\cap P_0]$. Then use

our (enhanced) *preservation* result.

Lemma. If a family $\mathcal{F}'\subset\omega^\omega$ is unbounded, then it remains so if we add any number of Cohen reals.

The Proof of the Lemma (uses nice names and factoring):

Proof. We take a nice name of a real $\dot h$ and have to show that it does not dominate $\mathcal{F}'$ in $V[G]$. Because $\dot h$ is a nice name of a countable set, there is a countable $J\subset I$ such that $\dot h$ is in fact a $\mathrm{Fn}(J,2)$-name. This means that $\dot h$ is a member of $V[G\cap\mathrm{Fn}(J,2)]$. The poset $\mathrm{Fn}(J,2)$ is "forcing isomorphic" to ${}^{<\omega}\omega$ and so we've already shown that $\mathcal{F}'$ is unbounded in $V[G\cap\mathrm{Fn}(J,2)]$.

This means that, in this inner model, $\dot h_G$ does not dominate $\mathcal{F}'$. Clearly, forcing or no forcing, this will not change when we move to the full $V[G]$.

Now to choose P_0 (reflection): We can assume that $\mathcal{F}=\{\dot f_\alpha:\alpha\in\omega_2\}$ and that $\{\dot r_\alpha:\alpha\in\omega_2\}$ enumerates all nice names for members of ω^ω.
Consider an arbitrary $\lambda<\omega_2$ and realize that there is an $\alpha_\lambda<\omega_2$ such that:

1. all nice $\mathrm{Fn}(\lambda,2)$-names in ω^ω occur among $\{\dot r_\alpha:\alpha<\alpha_\lambda\}$;
2. for each $\beta<\lambda$, $\dot f_\beta$ is a $\mathrm{Fn}(\alpha_\lambda,2)$-name;
3. for each $p\in\mathrm{Fn}(\lambda,2)$ and $\alpha<\lambda$, there is (this step is the key)
   a $q<p$ in $\mathrm{Fn}(\alpha_\lambda,2)$ and $\beta<\alpha_\lambda$ such that $q\Vdash\dot f_\beta\not<\dot r_\alpha$.

This is all just counting.

Now find an $\eta<\omega_2$ with cofinality ω_1 such that for all $\lambda<\eta$, α_λ is also less than η.

We'll find that, in $V[G_\eta]$:
$\{(\dot f_\alpha)_{G_\eta}:\alpha<\eta\}$ is unbounded. (basically from 3.)
By the preservation Lemma, it is unbounded in $V[G]$.

i.e. $\mathrm{Fn}(\omega_2,2)=\mathrm{Fn}(\lambda,2)*\mathrm{Fn}(\omega_2\setminus\lambda,2)$
preserves $\mathcal{F}'$ is unbounded

Proper Forcing

The notion of Proper Forcing is a brilliant creation of Shelah (with several important earlier contributions by others).

It is almost the minimum you can get away with by just saying that a poset must preserve ω_1. Recall the proof that ccc preserves ω_1: for each name τ of a function from ω into ORD, there's a ground model countable set F such that $\tau(n)\in F$ for each n.

'P is proper' can be formulated as [this is (incorrectly) simplified slightly]:
for each countable elementary submodel M (of V) with $P\in M$, there is a condition p (called (M,P)-generic) such that for each name $\tau\in M$ of an ordinal, $p\Vdash\tau\in M\cap\mathrm{ORD}$.

Because M is an elementary submodel, this is equivalent to: $D\cap M$ is predense below p for each dense $D\in M$. Of course $D\cap M$ is countable. It is something like

(M, P)-generic traps countably many maximal antichains to be essentially countable. You should think about why ccc posets and countably closed posets are proper.

Miller forcing is Proper. P is the set of all trees $T \subset \omega^{<\omega}$ such that each node $t \in T$ has either exactly one or infinitely many immediate successors and has an infinitely branching node somewhere above it. Remember the objects in the poset are the trees themselves, not the elements of the trees. We say that a tree T is stronger (or $<$) another $\tilde{T}$, if $T \subset \tilde{T}$. For a node $t \in T$, we let T_t denote the subtree consisting of those nodes which are compatible with t. Clearly $T_t < T$ and if $t \perp s$, then $T_t \perp T_s$.

P is definitely not ccc, but it is proper.

It is instructive to see properness proved:

Fix any countable M and let $\{D_n : n \in \omega\}$ enumerate the dense (open) subsets of P. We build a tree T which will be (M, P)-generic by a fusion (induction).

Fix any tree $T \in P \cap M$ (this is actually necessary) and find a $\tilde{T} \subset T$ which is (M, P)-generic.

We build a descending sequence $\{T_n : n \in \omega\}$ as follows. At stage n, we assume by induction that for each $t \in T_n$ such that t has no more than n branching predecessors, $(T_n)_t \in M$.

Notice that, L_n, the set of infinitely branching nodes of T_n which have exactly n infinitely branching predecessors, forms an antichain of T_n and $\{(T_n)_t : t \in L_n\}$ is an antichain of P.

For each $t \in L_n$ and each immediate successor $t^\frown \ell$ of t in T_n, choose a condition $S(t^\frown \ell) < (T_n)_{(t^\frown \ell)}$ which is in $D_n \cap M$ (argue using elementarity of M and density of D_n).

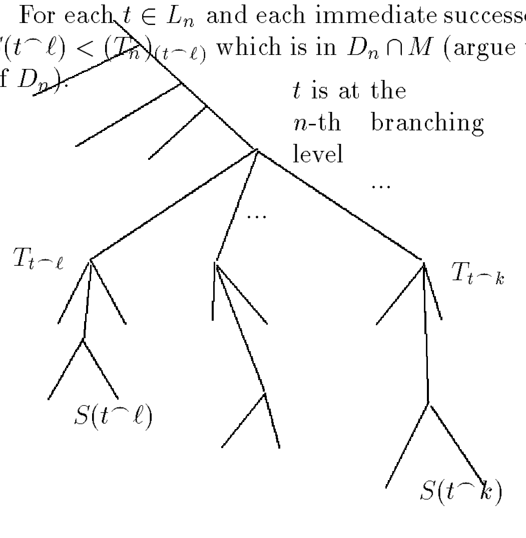

The tree $T_{n+1} = \bigcup\{S(t^\frown \ell) : t \in L_n, t^\frown \ell \in T_n\}$ is again a condition and each $t \in L_n$ remains a branching node in T_{n+1}. In fact, so too do all t in $\bigcup_{k<n} L_k$. (subtlety: $T_{n+1} \in M$ is not important but $S(t^\frown \ell) \in M$ is important)

The final tree is just $\tilde{T} = \bigcap_n T_n$ (the key idea is that we've ensured certain nodes (increasing finite set of them would suffice) remain branching).

We can see at a glance, that $\tilde{T}$ forces that $G \cap D_n$ contains one of the $S(t^\frown \ell)$ $(t \in L_n)$.

If P is a poset and $\dot{Q}$ is a P-name of a poset, then we set $P * Q$ equal to the set of all pairs $(p, \dot{q})$ with $p \in P$ and $q \in Q$.

Actually Q is a set of P-names together with a P-name of an order. If q_1 and q_2 are members of $\dot{Q}$, then we need various p in P to decide if $q_1 < q_2$ or $q_1 \perp q_2$.

The order on $P * Q$ is given by $(p_1, q_1) \leq (p_2, q_2)$ if $p_1 \leq p_2$ and $p_1 \Vdash q_1 \leq q_2$ and if G is P-generic, $V[P * Q] = V[P][Q_G]$.

For many posets at once: $\langle P_\alpha, Q_\alpha \rangle_{\alpha \in \kappa}$ can be the set of (partial) functions from κ into names such that, for each α, $p(\alpha) \in Q_\alpha$. If $p(\alpha)$ is not defined, then treat it as the 1 of Q_α. The ordering is given by $p_1 \leq p_0$ iff for each α, $p \restriction \alpha \Vdash_{P_\alpha} p_1(\alpha) \leq p_0(\alpha)$.

We do not have to take all such p. This makes a big difference in the behaviour. At a minimum, we take all functions with finite domain. If no others then this is called the finite support iteration. Another common example is the countable support iteration.

Major Theorems: finite support iteration of ccc is again ccc. (Δ-system Lemma; already $P * Q$ is ccc is subtle)

countable support iteration of countably closed is again countably closed.

countable support iteration of proper is again proper.

Let us look at Miller's result(s) that if we iterate Miller forcing (or Laver) then there are no q-points. The Miller model is now more noteworthy because Blass-Shelah proved that NCF holds in it.

An ultrafilter $\mathcal{U}$ on ω is a q-point, if for each increasing $h \in \omega^\omega$, there's a $U \in \mathcal{U}$ such that $|U \cap [h(n), h(n+1))| \leq 1$ for all n.

It is easy to see that, e.g., Martin's Axiom impies there are q-points. Also in the Cohen model and Random model they are known to exist.

Suppose that G is P_{ω_2}-generic (iterating Miller). In $V[G]$, let $\mathcal{U}$ be an ultrafilter on ω. Then there is a sequence, $\{\dot{U}_\alpha : \alpha \in \omega_2\}$, of names which constitute the members of $\mathcal{U}$ in $V[G]$.

Again, a nice name of a real has cardinality at most ω_1. Just as with Cohen forcing, there is a $\lambda < \omega_2$ with cofinality ω_1, such that $\{\dot{U}_\alpha, \omega \setminus \dot{U}_\alpha : \alpha \in \lambda\}$ is the list of "all" P_λ-names of reals (we must use properness of P_λ and $\mathrm{cf}(\lambda) = \omega_1$ for this).

We pass to the extension $V[G_\lambda]$. Clearly, $\{(\dot{U}_\alpha)_{G_\lambda} : \alpha < \lambda\}$ is an ultrafilter here.

I'll skip the verification that P_{ω_2}/P_λ is just iterated Miller forcing.

From the point of view of $V[G_\lambda]$ (now call it V): we have an ultrafilter $\mathcal{U}_\lambda$ (now call it $\mathcal{U}$) which might extend to a q-point after adding Miller reals. We show single stage does not.

Let g denote the first Miller real added and assume that τ is a Miller name of a real which meets $[g(n), g(n+1))$ in a singleton for each n. We will produce a condition T and a $U \in \mathcal{U}$ such that $T \Vdash \tau \cap U$ is empty.

For a branching node t of a condition T, note that

$|(\tau)_{T_t} \setminus \max t| \leq 1$ since $T_{t^\frown \ell} \Vdash |\tau \cap [\max t, \ell)| \leq 1$.

In a manner similar to the proof that the forcing is proper, we can also assume that $\forall \ell$, $F_{t^\frown \ell} = (\tau)_{T_{t^\frown \ell}} \setminus (\tau)_{T_t} \neq \emptyset$

Prune T so that if $s \neq t$, then $F_s \cap F_t$ is empty.

Inductively partition T into two conditions. One of them works since $\tau_T^+ = \bigcup\{F_t : t \in T'\}$ for $T' \leq T$.

Contrast with Cohen: Suppose $g \in \omega^\omega$ is any real and let h be added by Cohen forcing P_C. Define W to be all those $h(k)$'s which are the only ones in their $[g(n), g(n+1))$-interval.

For each infinite $U \subset \omega$,

$D_U = \{p \in P_C : p \Vdash U \cap W \neq \emptyset\}$ is dense.

Indeed choose any n large enough so that

$U \cap [g(n), g(n+1))$ is not empty, and

$p(j) < g(n)$ for all $j \in k = \mathrm{dom}(p)$. Extend p to $q = p \cup \{(k, i)\} \cup \{(k+1, g(n+1))\}$.

Then q is in D_U.

Another example: Compact separable radial spaces are Frèchet.

We analyze the problem and look for vulnerability:

radial + not Frèchet implies

ω is dense and there's a converging free ω_1-sequence $\{x_\alpha : \alpha \leq \omega_1\}$

(size $\leq \omega_1$ here is quite important)

By radial, there are $Y_\alpha \subset \omega$ converging to x_α. Also, using some dominating functions, an increasing mod finite sequence $\{A_\alpha : \alpha \in \omega_1\}$ of subsets of ω such that $Y_\alpha \subset^* A_\alpha$ and the closure of A_α is disjoint from $[x_{\alpha+1}, x_{\omega_1}]$.

Force a $Y \subset \omega$ such that

1) $|Y \cap Y_\alpha|$ is infinite for each $\alpha < \omega_1$

2) for each $Z \in [Y]^\omega$, $|Z \cap A_\alpha| = \omega$ for some α. (old? new?)

Well, this means that x_{ω_1} is in $\overline{Y}$ and yet no $Z \subset Y$ converges to x_{ω_1} (radial is destroyed).

But: there are other examples which must be killed (will doing so introduce a bad Z to this one?=preservation)
and as we kill them off, we keep introducing more. (kill a subspace might work =reflection)

We know in advance that we design $P_{\omega_2} = \langle P_\alpha, \dot{Q}_\alpha \rangle_{\alpha \in \omega_2}$ (finite support).

In $V[G_{\omega_2}]$, consider a space with $\{\dot{Y}_\alpha : \alpha \in \omega_1\}$, $\{\dot{W}_\alpha : \alpha \in \omega_1\}$, W_α regular closed so that $W_\alpha \cap [x_0, x_\alpha]$ is empty and $[x_{\alpha+1}, x_{\omega_1}] \subset \text{int } W_\alpha$. ω is dense, so $W_\alpha \cap \omega$ is all we need.

As before, there's a $\lambda < \omega_2$ so that $V[G_\lambda] \models Y_\alpha, W_\alpha, x_\alpha$ have all the above relationships.

So Q_λ (or some Q_γ $(\gamma \geq \lambda)$) should introduce Y and the A_α's as described before (i.e., $|A_\alpha \cap W_\alpha|$ should be finite).

Then, $V[G_{\lambda+1}] \models |Y \cap Y_\alpha| = \omega$ and

(*) for all $Z \in [Y]^\omega$, $Z \cap A_\alpha$ is infinite for some α.

So, (*) must be preserved. We mentioned that $P_{\mathcal{F}}$ does this. So hard work needed to design Q_λ to all "preserve towers". Finally

Lemma: finite support iteration of ccc tower preserving, also preserves towers.

Set $\mathcal{A} = \{A_\alpha : \alpha \in \omega_1\}$ an $\subset^*$-increasing maximal chain in $Y \subset \omega$. Suppose that $\langle P_\beta, Q_\beta \rangle$ is a FS ccc iteration and that τ is a nice name of a subset of Y.

Let δ be minimal such that $\tau \subset P_\delta \times Y$ (proceed by induction on δ).

1) If $\delta = \gamma + 1$. Then $\Vdash_{P_\gamma} \mathcal{A}$ is maximal. Pass to $V[G_\gamma] \models Q_\gamma$ preserves $\mathcal{A}$ is maximal, hence $V[G_\gamma][G_{\gamma+1}/G_\gamma] \models \mathcal{A}$ is a maximal, i.e., $V[G_\delta] \models \mathcal{A}$ is maximal.

2) if $\text{cf}(\delta) = \omega$, then fix $\delta_m : m \in \omega$ increasing cofinal in δ.
For each $\alpha \in \omega_1$, assume there is a p_α and an n_α such that $\tau \cap A_\alpha \subset n_\alpha$. Also let m_α be such that $p_\alpha \in P_{\delta_{m_\alpha}}$.

So, there is an uncountable $I \subset \omega_1$ and a pair n, m such that $n_\alpha = n$ and $m_\alpha = m$ for all $\alpha \in I$.

Special fact about ccc poset P_{δ_m}: there is a generic filter G such that $G \cap \{p_\alpha : \alpha \in I\}$ is uncountable. Fix such a G. Pass to the extension $V[G]$.

In $V[G]$, let $J = \{\alpha \in I : p_\alpha \in G\}$. Also consider the set

$$\tau^+ = \{\ell \in Y : (\exists p \in P_\delta)\ p \restriction \delta_m \in G \text{ and } p \Vdash \ell \in \tau\}$$

i.e., those ℓ that still have a chance to get into τ. Fix an $\alpha \in J$ such that $A_\alpha \cap \tau^+$ is infinite. Then pick any $\ell \in A_\alpha \cap \tau^+$ such that $\ell \geq n$. So there is a p such that $p \restriction \delta_m \in G$ and $p \Vdash \ell \in \tau$.

But now, p_α and p are compatible because p_α and $p \restriction \delta_m$ are both in G and so are compatible. Let q be any common extension.

$q \Vdash \tau \cap A_\alpha \not\subset n_\alpha$ - a contradiction to $q \leq p_\alpha$