Diploma Thesis
on
The Modal Logic of Forcing

by
Colin Jakob Rittberg

September 14, 2010

Thanks

First and most important of all I would like to thank Prof. Dr. Ralf Schindler. He has been very influential throughout my whole career at the University of Münster, guiding me and helping me on more than one occasion. He also introduced me to the interesting field of mathematical logic and provided me with this subject for my diploma thesis. He had the fundamental idea for the buttons constructed in this work and helped whenever I needed help.

Secondly I would like to thank Prof. Dr. Joel David Hamkins and Prof. Dr. Benedikt Löwe for their interesting paper and for all the help they were always willing to give throughout my work on this thesis.

I also want to thank Dr. Philipp Doebler and Dominik Adolf. Both have been of great help throughout my studies and especially Dominik Adolf helped me a great deal with some of the calculations and the so-called "dirty-work" in this thesis.

And I want to thank Max Giersberg who wrote his thesis at about the same time as I did and who was always willing to listen to my problems. It was great to have someone who understood the mathematics without always knowing the answers to my questions.

Last but not least I would like to thank my mother, for she was the one who encouraged me and cared for me whenever I got stuck or big mistakes rocked my boat.

Contents

Chapter 0

Abstract

This thesis is based on the paper "The Modal Logic of Forcing"[5] by Prof. Dr. J.D. Hamkins and Prof. Dr. B. Löwe. We will follow their lead and interpret the modal operators $\Diamond$ and $\Box$ in the language of set theory such that $\Diamond\phi$ is taken to mean "there is a forcing extension in which ϕ holds" and $\Box\phi$ abbreviates "ϕ holds in every forcing-extension". The above mentioned paper then detaills how to link forcing and modal logic together and, modifying the arguments in [5] slightly, enables to prove that the modal logic of forcing is S4.2. We then take a step further and restrict ourselves to ω-closed forcing, obtaining the main result that the modal logic of ω-closed forcing is also S4.2. A simple generalisation proves the same for $< \kappa$-closed forcings where κ is definable by some formula ϕ (see chapter 4 for details).

While the reader is expected to understand all the basic concepts of forcing as presented in [10], in chapter 1 we see a thorough introduction to modal logic from the very basics to all the results needed in this thesis. In chapter 2 a rough outline of our argumentation is presented, notations are given and some basic facts not covered by [10] are proven. Then we define countably many sentences in $\mathcal{L}_\in$ which will be used as the backbone of our argument in the following chapters. Chapter 3 is basically a reformulation of the first part of [5] while chapter 4 takes the ideas presented in chapter 3 and applies them to special classes of forcings.

Chapter 1

Modal Logic

This chapter follows the structure of the excellent book "Modal Logic" by P. Blackburn, M. de Rijke and Y. Venema[1] very closely.
The language of modal logic[1] is defined by enriching the language of propositional logic with a new symbol: $\Diamond$. The interpretation of $\Diamond$ makes it a very useful tool in expressing relations between structures. And that is all modal logic is about: relational structures. These strucures are called *Kripke frames*.
As always with logic, truth is a key concept and the modal language gives rise to different kinds of truth: local and global, and an even more general concept of truth: validity. Local and global truth need a fixed valuation which tells us whether or not some statement ϕ is true at a given point in our Kripke Frame or not. Local truth means truth in one point, global truth means truth in all points of the given structure. Validity on the other hand means "global truth for all valuations over a given frame".
Once we define this we have all we need to state what is meant by "a *modal logic*" and we will then quickly focus on well-behaved ones: the *normal modal logics*. They are essentially sets of formulas with nice closure conditions. There is a minimal normal modal logic, K, and by adding certain new formulas to K one generates new normal modal logics. One of them will be of particular importance in this thesis: S4.2.
Thus far the normal modal logics have been semantically defined. On the other hand a class of frames $\mathbb{F}$ gives rise to a completely syntactically defined logic: $\Lambda_{\mathbb{F}}$, the set of formulas valid on every frame in $\mathbb{F}$. It seems natural to ask under what circumstances these logics coincide.
Section 2 of this chapter discusses the basic notations of *soundness* and *completeness* needed to answer this question. Soundness means that for a semantically defined normal modal logic Λ every $\phi \in \Lambda$ is true on the class of frames $\mathbb{F}$, which is just $\Lambda \subseteq \Lambda_{\mathbb{F}}$. Completeness implies that the semantically defined logic completely captures all validities in the class of frames. Therefore showing

[1] Just like the expression "the first order logic" it is a slight abuse of notation to talk about the language of modal logic. There are many of them! But we will not bother with this nicety, mainly because we will only need the so-called "basic modal logic".

soundness and completeness of some logic Λ with respect to a class of frames establishes an equality of Λ and the logic that is syntactically defined by the class of frames.
As will be seen, showing soundness boils down to an easy task. Section 3 then discusses one method to show completeness: canonical models. Generated from maximal consistent sets of formulas, canonical models have very nice properties that will lead directly to the Canonical Model Theorem 1.3.7, which states that any normal modal logic is (strongly) complete with respect to its canonical model. This will help us to show that the above mentioned logics K, S4, S4.2 and S5 are complete for certain classes of frames.
S4.2 suffices for the discussion of the modal logic of forcing in this thesis, while S4 will be a lower bound for the modal logic of c.c.c. forcings. S5 is the modal logic of forcing restricted to a model where every button is pushed (see definitions in chaper 2). This will not be proven in this thesis and we refer the reader to Hamkins paper [4] for a proof. Other logics may occur when one restricts oneself to another class of forcings. For example is a work by Hamkins and Löwe to appear in which they prove the modal logic of $\mathrm{Col}(\omega, \theta)$-forcing to be S4.3, a normal modal logic not discussed in this thesis (see [6]).
Once the completeness results are established we will see a surprising result: every formula not in S4.2 can be falsified on a finite Kripke Model. This property is called the Finite Frame Property and we will introduce the filtration method in order to show it. We will make heavy use of the finite frame property when we show that the modal logic of forcing is S4.2.

1.1 The Basics

As stated, the modal language, and hereby I mean the basic modal language, is defined by enriching the propositional language by a modal operator.

Definition 1.1.1 The *basic modal language* is defined by a countable set of propositional letters Φ and an unary modal operator $\Diamond$. The well-defined formulas of the basic modal language are defined by recursion as follows:

- every p is a formula for every $p \in \Phi$.
- $\bot$ is a formula.
- if ϕ is a formula, then $\neg\phi$ is a formula.
- if ϕ and ψ are formulas, then $\phi \vee \psi$ is a formula.
- if ϕ is a formula, then $\Diamond\phi$ is a formula.

Just as with the quantifiers in first order logic we have a duality with the modal operator. The $\Box$ is then interpreted as shorthand for:

$$\Box\phi := \neg\Diamond\neg\phi$$

This duality actually only holds for modal logics which contain the (Dual) axiom, as we will see later. But throughout this paper we will only be interested in so called *normal modal logics*, and those contain (Dual) by definition. See Definition 1.1.8.

The $\Diamond$ is called diamond and $\Box$ is called box. They can be thought of as "something is *possible*" and "something is *necessary*" respectively. But as the modal language gives rise to a variety of different interpretations, the possible/necessary interpretation is not always the best. There are two more readings of the modal operators that have been extremely influential:
First the epistemic reading "$\Box\phi$= the agent knows ϕ". Here it is normal to write Kϕ for $\Box\phi$. The diamond is then interpreted as "it is consistent with what the agent knows" and is usually called M.
Secondly, in provability logic $\Box\phi$ is read as "ϕ is provable".
But there are even more readings. In the forcing interpretation of the modal language, for example, the diamond is thought of as "there is a forcing extension such that...".
From all those different readings comes the great power of the modal language. But the power can be extended still:

Remark 1.1.2 One does not have to restrict oneself to a modal language with only one modal operator or the demand on the operator(s) to take only a single formula as argument. But we will not need this generalization, so we will only consider the so-called *basic modal language*, which was defined above. We will also refer to this language as *the* modal language.

With this generalized modal language one can obviously read the (now many different) modal operators as even more different statements.
Now that the modal language is defined, we are interested in a definition of truth. One defines Kripke Frames to fix the relational structure. Enriching these Kripke Frames by a valuation generates Kripke Models and allows a truth definition.

Definition 1.1.3 A *Kripke Frame* for the modal language is a pair $\mathcal{F} = (\mathcal{W}, \mathcal{R})$ such that:

(i) $\mathcal{W}$ is a non-empty set

(ii) $\mathcal{R}$ is a binary relation on $\mathcal{W}$

The elements of $\mathcal{W}$ are called nodes, states, or worlds. $\mathcal{R}$ is referred to as the accessibility relation.
We also say *frame* when we mean *Kripke* frame.
A *Kripke Model* of the modal language is a pair $\mathcal{M} = (\mathcal{F}, \mathcal{V})$, where $\mathcal{F}$ is a frame for the modal logic and $\mathcal{V}$ is a function

$$\mathcal{V} : \Phi \to \mathscr{P}(\mathcal{W})$$

$\mathcal{V}$ is called the valuation and for $\mathcal{M} = (\mathcal{F}, \mathcal{V})$ we say that $\mathcal{M}$ is based on $\mathcal{F}$ or $\mathcal{F}$ is the underlying frame of $\mathcal{M}$.
Again we omit *Kripke* and say *model* most of the time.

With the above notation we will sometimes write $w \in \mathcal{F}$ or "$\mathcal{F}$ has property P" when in fact we mean $w \in \mathcal{W}$ or "$\mathcal{R}$ has property P". The same goes for models.

We defined $\mathcal{V}$ in such a way that we can construct a definition of truth where p is true at some world w iff $w \in \mathcal{V}(p)$. We can think of this as the world w having the property p.

Definition 1.1.4 For $\mathcal{M}$=($\mathcal{W}$,$\mathcal{R}$,$\mathcal{V}$) and $w \in \mathcal{W}$ we define the notion of a formula ϕ being satisfied (or true) in $\mathcal{M}$ at a state w inductively:

$\mathcal{M}, w \models p$ iff $w \in \mathcal{V}(p)$, where $p \in \Phi$

$\mathcal{M}, w \models \perp$ iff never

$\mathcal{M}, w \models \neg\phi$ iff not $\mathcal{M}, w \models \phi$

$\mathcal{M}, w \models \phi \vee \psi$ iff $\mathcal{M}, w \models \phi$ or $\mathcal{M}, w \models \psi$

$\mathcal{M}, w \models \Diamond\phi$ iff there is some $v \in \mathcal{W}$ such that $w\mathcal{R}v$ and $\mathcal{M}, v \models \phi$

As always, a set of formulas Σ is true at a world w in a model $\mathcal{M}$ if all members of Σ are true at w. This is abbreviated by $\mathcal{M}, w \models \Sigma$

Now a formula is *locally true* if it holds at a world in $\mathcal{W}$. A formula is *globally true* if it holds in every world of $\mathcal{W}$.
Consider now a reflexive frame, i.e. a frame whose relation is reflexive. Then $w\mathcal{R}w$ for all $w \in \mathcal{W}$. So for every valuation: if $\mathcal{M}, w \models \phi$ then $\mathcal{M}, w \models \Diamond\phi$. Hence $\mathcal{M}, w \models \phi \rightarrow \Diamond\phi$ for all valuations and all worlds w. So we have found a formula that is true for all frames with a reflexive relation. Those formulas are called valid on the class of reflexive frames.

Definition 1.1.5 For a formula ϕ we say:

ϕ is valid at a world w in a frame $\mathcal{F}$, abbreviated by $\mathcal{F}, w \models \phi$, if ϕ is true at w for every model $(\mathcal{F}, \mathcal{V})$ based on $\mathcal{F}$.

ϕ is *valid in a frame $\mathcal{F}$*, abbreviated by $\mathcal{F} \models \phi$, if ϕ is valid at every state in $\mathcal{F}$.

ϕ is *valid on a class of frames* $\mathbb{F}$, $\mathbb{F} \models \phi$, if ϕ is valid on every frame $\mathcal{F}$ in $\mathbb{F}$.

ϕ is valid, $\models \phi$, if ϕ is valid on the class of all frames.

The set of all formulas valid in a class of frames $\mathbb{F}$ is referred to as the logic of $\mathbb{F}$ and is called $\Lambda_{\mathbb{F}}$

The following example points out the difference between truth and validity.

Example 1.1.6 As mentioned above, truth is a local attribute while validity is global. This difference becomes clear when one thinks of the formula $\phi \vee \psi$. If, for some model $\mathcal{M}$=($\mathcal{F}$,$\mathcal{V}$) and some world w, $\mathcal{M}, w \models \phi \vee \psi$ then ϕ or ψ (or both) are true in $\mathcal{M}$ at w. On the other hand, if $\mathcal{F} \models \phi \vee \psi$, it is generally not true that $\mathcal{F} \models \phi$ or $\mathcal{F} \models \psi$. $p \vee \neg p$ gives a simple counter-example.

We now know what our language looks like, what frames and models are, and what is meant by truth and validity. We have also seen the syntactically defined logic for a class of frames: $\Lambda_{\mathbb{F}}$. The next step is to define semantically what is meant by *logic* and what is considered well-behaved. This gives a good understanding of the logic, but also obliges one to link those two definitions of logic. As said in the introduction of this chapter, section 2 will provide us with the definitions and section 3 with the tools needed to complete this task.

Definition 1.1.7 A set Λ of modal formulas that contains all propositional tautologies and is closed under modus ponens (i.e., $\phi \in \Lambda$ and $\phi \to \psi \in \Lambda$ implies $\psi \in \Lambda$) and uniform substitution (i.e., if $\psi \in \Lambda$ then all substitution instances of ψ belong to Λ) is called a *modal logic*.

Elements of Λ are called *theorems* and we write $\vdash_\Lambda \phi$ if ϕ is a theorem of Λ and $\nvdash_\Lambda \phi$ if not.

As a convention, the "modal" shall usually be droped and we only talk about logics.
Some logics seem rather odd. Imagine for example a logic Λ that contains some formula ϕ, but not $\Box\phi$. To build a frame that satisfies this logic one has to build a world w and a valuation such that $\Box\phi$ does not hold at w, even though ϕ holds in all worlds. Such w are called non-normal worlds, for if $\Box\phi$ does not hold for all $\phi \in \Lambda$, these worlds see other worlds where even tautologies are wrong[2]. To navigate around this one defines:

Definition 1.1.8 A (modal) logic Λ is *normal* if it contains the formulas

(K) $\Box(p \to q) \to (\Box p \to \Box q)$,

(Dual) $\Diamond p \leftrightarrow \neg\Box\neg p$,

[2]A model for a non-normal modal logic is of the form $\mathcal{M} = (\mathcal{W}, \mathcal{N}, \mathcal{R}, \mathcal{V})$, where $\mathcal{W}$ and $\mathcal{R}$ are as above and $\mathcal{N} \subseteq \mathcal{W}$. The worlds in $\mathcal{N}$ are called normal and worlds in $\mathcal{W} \setminus \mathcal{N}$ non-normal. The truth conditions for the truth functions $\vee, \wedge, \neg$ are as in definition 1.1.4. The truth conditions for $\Box$ and $\Diamond$ at normal worlds are also as in 1.1.4, but if w is a non-normal world one defines:

- $\mathcal{M}, w \models \Box\phi$ iff never
- $\mathcal{M}, w \models \Diamond\phi$ iff always

Then the tautologie $\Box(p \vee \neg p)$ fails at w. As we are not intested in non-normal logics in this thesis we refer the reader to [9] chapter 4 for a thorough treatment.

and is closed under necessitation:

(nec) if $\vdash_\Lambda \phi$ then $\vdash_\Lambda \Box\phi$.

With this definition our difficulty disappears by demanding that Λ has to be normal. When we say logic in this thesis we usually mean *normal* modal logic. We are now ready to define the important concept of consistency. It will enable us to define *canonical models* in section 3.

Definition 1.1.9 Let $\Gamma \cup \{\phi\}$ be a set of formulas and Λ be a logic. We say that ϕ *is deducible in* Λ *from* Γ or ϕ *is* Λ*-deducible from* Γ if $\vdash_\Lambda \phi$ or there are formulas $\psi_1, ..., \psi_n \in \Gamma$ such that

$$\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \phi.$$

If this is the case we write $\Gamma \vdash_\Lambda \phi$ and $\Gamma \nvdash_\Lambda \phi$ otherwise.
A set of formulas Γ is Λ-*consistent* if $\Gamma \nvdash_\Lambda \bot$, and Λ-*inconsistent* otherwise.
Consistency of ϕ means that $\{\phi\}$ is Λ-consistent.

1.2 Soundness and Completeness

We are looking for a class of frames $\mathbb{F}$ such that Λ=$\Lambda_\mathbb{F}$ for a given logic Λ. The direction $\Lambda \subseteq \Lambda_\mathbb{F}$ is referred to as *soundness* and the other direction as *completeness*. Since it is far more difficult for a given class of frames to show completeness than soundness, we begin gradually and start with the soundness definition.

Definition 1.2.1 Let $\mathbb{F}$ be a class of structures[3]. A normal logic Λ is *sound with respect to* $\mathbb{F}$ if $\Lambda \subseteq \Lambda_\mathbb{F}$

A simple corollary eases the use of this definition.

Lemma 1.2.2 *Λ is sound w.r.t.*[4] *$\mathbb{F}$ iff for all formulas ϕ and all $\mathcal{F} \in \mathbb{F}$: $\vdash_\Lambda \phi$ implies $\mathcal{F} \models \phi$.*

Proof. Let $\Lambda \subseteq \Lambda_\mathbb{F}$ and let $\vdash_\Lambda \phi$. Then, by definition, $\phi \in \Lambda \subseteq \Lambda_\mathbb{F}$. So ϕ holds on all structures in $\mathbb{F}$, i.e. $\mathcal{F} \models \phi \ \forall \mathcal{F} \in \mathbb{F}$.

On the other hand: if there is a $\phi \in \Lambda$ with $\neg(\phi \in \Lambda_\mathbb{F})$, then there is a $\mathcal{F} \in \mathbb{F}$ with $\mathcal{F} \not\models \phi$, i.e. $\vdash_\Lambda \phi$ but $\mathcal{F} \not\models \phi$. □

We now have established the left to right inclusion from our main task. To apply this to the cases that remain interesting throughout this thesis, K, S4, S4.2 and S5, we have to define them first:

Consider the following axioms:

[3]We refer to structures when we want to talk about models as well as frames.
[4]means: with respect to

(K) $\Box(p \to q) \to (\Box p \to \Box q)$

(Dual) $\Diamond p \leftrightarrow \neg\Box\neg p$

(S) $\Box p \to p$

(4) $\Box p \to \Box\Box p$

(.2) $\Diamond\Box p \to \Box\Diamond p$

(5) $\Diamond\Box p \to p$

One then defines the logic K(see[5]) to be the minimal normal modal logic, i.e., K consists of all propositional tautologies, the axioms (K) and (Dual) and is closed under modus ponens, uniform substitution, and necessitation.

Definition 1.2.3 For axioms $A_1, ..., A_n$ let $KA_1...A_n$ be the normal logic generated by $A_1, ..., A_n$. One then defines:

$$S4 = KS4$$
$$S4.2 = KS4.2$$
$$S5 = KS45$$

There are (countably) many more axioms and named logics one might consider. See [1] for a further discussion.

Lemma 1.2.4 *The above defined modal logics we achieve the following soundness results:*

(i) K is sound with respect to the class of all frames.

(ii) S4 is sound w.r.t. the class of reflexive and transitive frames.

(iii) S4.2 is sound w.r.t. the class of frames whose relation is a directed partial order[6]*.*

(iv) S5 is sound w.r.t. the class whose relation is an equivalence relation.

Proof. (i) Let $\mathcal{F}$=($\mathcal{W}$,$\mathcal{R}$) be an arbitrary frame. To show that modus ponens holds, let $\mathcal{F} \models \phi \to \psi$ and $\mathcal{F} \models \phi$. Then for all $w \in \mathcal{W}$ and all valuations $\mathcal{V}$ $(\mathcal{F}, \mathcal{V}), w \models \phi \to \psi$ and $(\mathcal{F}, \mathcal{V}), w \models \phi$. Because ($\mathcal{F}$,$\mathcal{V}$) is a model we have $(\mathcal{F}, \mathcal{V}), w \models (\phi \to \psi) \wedge \phi$ by definition. So $(\mathcal{F}, \mathcal{V}), w \models \psi$, and by arbitrarity modus ponens holds.
If $\mathcal{F} \models p$, then for all $w \in \mathcal{W}$ $\mathcal{F}, w \models p$. So for all v with $w\mathcal{R}v$ $\mathcal{F}, v \models p$

[5]Note the difference between the logic K and the axiom (K).

[6]A partial order R is called *directed* iff $\forall x, y, z(xRy \wedge xRz \to \exists u(yRu \wedge zRu))$.

(by $\mathcal{F} \models p$), so $\mathcal{F}, w \models \Box p$. (nec) follows from arbitrarity.
To show that (K) holds in $\mathcal{F}$, let $\mathcal{F} \models \Box(p \to q) \wedge \Box p$. Then at every $w \in \mathcal{W}$ $\mathcal{F}, w \models \Box(p \to q) \wedge \Box p$. So for every v with $w\mathcal{R}v$ $\mathcal{F}, v \models (p \to q) \wedge p$. By modus pones $\mathcal{F}, v \models q$ for all $w\mathcal{R}v$, hence $\mathcal{F} \models (K)$.
For (Dual) consider $\mathcal{F} \models \Diamond p$. Then for all $w \in \mathcal{W}$ $\mathcal{F}, w \models \Diamond p$. Now suppose there is a $v \in \mathcal{W}$ with $\mathcal{F}, v \models \Box\neg p$. Then for all v' with $v\mathcal{R}v'$ $\mathcal{F}, v' \models \neg p$. But by $\mathcal{F} \models \Diamond p$ there is a v' with $v\mathcal{R}v'$ and $\mathcal{F}, v' \models p$. Contradiction!
On the other hand, $\mathcal{F} \models \neg\Box\neg p$ means that for every $w \in \mathcal{W}$ there is a $v \in \mathcal{W}$ such that wRv and $\mathcal{F}, v \models p$. Hence $\mathcal{F} \models \Diamond p$ as desired.
Now let $\phi(p)$ be a formula with at least one occurence of the propositional variable p. Write $\phi(p \mid q)$ for the formula obtained by taking ϕ and replacing every instance of p by q, where q is again a propositional variable[7]. If $\mathcal{F} \models \phi(p)$, then $(\mathcal{F}, \mathcal{V}) \models \phi(p)$ for every valuation $\mathcal{V}$. But that means that $(\mathcal{F}, \mathcal{V}) \models \phi(p \mid q)$ and hence $\mathcal{F} \models \phi(p \mid q)$.
Because models cannot falsify tautologies, the above argument also works for tautologies.
This completes the proof of (i).

(ii) Because (K) holds on arbitrary frames it only remains to show the axioms (S) and (4).
For (4) it suffices to show that $\Diamond\Diamond p \to \Diamond p$ holds[8].
So let $\mathcal{F} \models \Diamond\Diamond p$ where $\mathcal{F}$=($\mathcal{W}$,$\mathcal{R}$) such that $\mathcal{R}$ is reflexive and transitive. Then $\mathcal{F}, w \models \Diamond\Diamond p$ for all $w \in \mathcal{W}$. Let $v, v' \in \mathcal{W}$ be such that $w\mathcal{R}v$ and $\mathcal{F}, v \models \Diamond p$ and $v\mathcal{R}v'$ and $\mathcal{F}, v' \models p$. By transitivity of $\mathcal{R}$ it follows that $w\mathcal{R}v'$. Hence $\mathcal{F}, w \models \Diamond p$.
For (S)
Let $\mathcal{F} \models \Box p$. So for all $w \in \mathcal{W}$ we have $\mathcal{F}, w \models \Box p$. $w\mathcal{R}w$ because $\mathcal{R}$ is reflexive, hence $\mathcal{F}, w \models p$ for all $w \in \mathcal{W}$. So $\mathcal{F} \models p$.

(iii) Let $\mathcal{F}$=($\mathcal{W}$,$\mathcal{R}$) be a frame and $\mathcal{R}$ be a directed partial order. Since $\mathcal{R}$ is reflexive and transitive only (.2) remains to be shown.
Let $\mathcal{F} \models \Diamond\Box p$, i.e. for all worlds w_0 there is a world w_1 such that w_0Rw_1, $\mathcal{F}, w_0 \models \Diamond\Box p$ and $\mathcal{F}, w_1 \models \Box p$. Let w_2 be arbitrary with w_0Rw_2. By directedness there is a w_3 with $w_1Rw_3 \wedge w_2Rw_3$. So $\mathcal{F}, w_3 \models p$ which implies $\mathcal{F}, w_2 \models \Diamond p$ and because w_2 was arbitrary $\mathcal{F}, w_0 \models \Box\Diamond p$.

(iv) Let $\mathcal{F}$=($\mathcal{W}$,$\mathcal{R}$) where $\mathcal{R}$ is an equivalence relation. It only remains to show (5) by (ii).
Let $\mathcal{F} \models \Diamond\Box p$. So for all $w \in \mathcal{W}$ there is a $v \in \mathcal{W}$ with $w\mathcal{R}v$ and

[7] This is called the uniform substitution of p by q in ϕ.
[8] This is because of the equivalence of the formula with (S) in K:
$\Diamond\Diamond p \to \Diamond p$ iff (Dual) $\neg(\neg\Box\neg\neg\Box\neg p) \vee \neg\Box\neg p$ iff $\Box\Box\neg p \vee \neg\Box\neg p$ iff $\Box\neg p \to \Box\Box\neg p$ iff (uni sub) $\Box p \to \Box\Box p$
The formula $\Diamond\Diamond p \to \Diamond p$ is sometimes also called (T).

$\mathcal{F}, v \models \Box p$. So for all $v' \in \mathcal{W}$ with $v\mathcal{R}v'$ we have $\mathcal{F}, v' \models p$. But $\mathcal{R}$ is symmetric, so $v\mathcal{R}w$ and therefore $\mathcal{F}, w \models p$

□

The task of showing that the logic generated by a given class of frames is completely captured by a (semantically defined) logic leads to the concept of completeness. There are two different kinds of completeness: strong and weak. As expected, strong completeness implies weak completeness. We will make use of this, for canonical models show strong completeness. We then conclude weak completeness and thereby complete our task.

Definition 1.2.5 Let $\mathbb{F}$ be a class of structures and $\Gamma \cup \{\phi\}$ be a set of fomulas. We say that ϕ is a *local semantic consequence of* Γ *over* $\mathbb{F}$, abbreviated by $\Gamma \models_{\mathbb{F}} \phi$, if for all models $\mathcal{M}$ built out of $\mathbb{F}$ (see [9]) and all states w in $\mathcal{M}$:
if $\mathcal{M}, w \models \Gamma$ then $\mathcal{M}, w \models \phi$.
A logic Λ is *strongly complete w.r.t.* $\mathbb{F}$ if $\Gamma \models_{\mathbb{F}} \phi$ implies $\Gamma \vdash_{\Lambda} \phi$.
A logic Λ is *weakly complete w.r.t.* $\mathbb{F}$ if for any formula ϕ: if $\mathbb{F} \models \phi$ then $\vdash_{\Lambda} \phi$.

The following corollary states again what already has been said.

Corollary 1.2.6 Λ *strongly complete w.r.t.* $\mathbb{F} \Rightarrow \Lambda$ *weakly complete w.r.t.* $\mathbb{F}$

Proof. Take $\Gamma = \emptyset$. By strong completeness $\emptyset \models_{\mathbb{F}} \phi$ implies $\vdash_{\Lambda} \phi$. But $\emptyset \models_{\mathbb{F}} \phi$ is just $\mathbb{F} \models \phi$. □

The following corollary just says that what we have defined is precisely what is needed to complete our task.

Corollary 1.2.7 Λ *is weakly complete w.r.t.* $\mathbb{F}$ *iff* $\Lambda_{\mathbb{F}} \subseteq \Lambda$.

Proof. "$\Rightarrow$" Let $\phi \in \Lambda_{\mathbb{F}}$. Then $\mathbb{F} \models \phi$ by definition of $\Lambda_{\mathbb{F}}$. So by weak completeness $\vdash_{\Lambda} \phi$ i.e. $\phi \in \Lambda$.
"$\Leftarrow$" Let $\mathbb{F} \models \phi$, then $\phi \in \Lambda_{\mathbb{F}} \subseteq \Lambda$, i.e. $\phi \in \Lambda$. □

Now that we have defined soundness and completeness, we have all we need to link syntactically and semantically specified logics but let us first have a few words on the symmetry of this definition to the one given in [5]:
Hamkins and Löwe define a class of frames $\mathbb{F}$ to be complete for a logic Λ if every $\mathcal{F} \in \mathbb{F}$ is a Λ-frame and any ϕ true in all Kripke Models having frames in $\mathbb{F}$ is in Λ.
Here a Λ-frame is a frame such that every model based on this frame satisfies Λ.
This definition is essentially dual to the one given in this thesis in the following way:

[9] i.e., if $\mathbb{F}$ is a class of frames, $\mathcal{M}$ has a frame in $\mathbb{F}$, if $\mathbb{F}$ is a class of models, $\mathcal{M}$ is an element of $\mathbb{F}$.

Saying that $\mathbb{F}$ is a class of Λ-frames just means that Λ is sound w.r.t. $\mathbb{F}$, while the second condition can be restated as: if $\mathbb{F} \models \phi$ then $\vdash_\Lambda \phi$, which is precicely the definition of weak completeness of Λ w.r.t. $\mathbb{F}$.
The difference is of course that Hamkins and Löwe defined what it means for a class of *frames* to be complete for some logic, while we defined what it means for some *logic* to be sound and complete for some class of frames.

We will need a technical lemma:

Lemma 1.2.8 (Contradiction Lemma) *Let $\Gamma \cup \{\phi\}$ be a set of formulas such that $\Gamma \cup \{\phi\} \vdash_\Lambda \bot$. Then $\Gamma \vdash_\Lambda \neg\phi$.*

Proof. A proof can be found for example in [10], where the proof is done for first order logic, but it also works for modal formulas. □

The following lemma gives us a neat re-definition of strong completeness.

Lemma 1.2.9 *A logic Λ is strongly complete w.r.t. a class of structures $\mathbb{F}$ iff every Λ-consistent set of formulas Γ is satisfiable on some $\mathcal{F} \in \mathbb{F}$, i.e., there is a state w in $\mathcal{F}$ such that every $\phi \in \Gamma$ is true at w in $\mathcal{F}$.*

Proof. To show right to left argue by contraposition:
Suppose Λ is not strongly complete w.r.t. $\mathbb{F}$. Then there is a set of formulas $\Gamma \cup \{\phi\}$ such that $\Gamma \vdash_{\mathbb{F}} \phi$ but $\Gamma \nvdash_\Lambda \phi$. Then $\Gamma \cup \{\neg\phi\}$ is Λ-consistent by the Contradiction Lemma 1.2.8, but not satisfiable on any structure in $\mathbb{F}$.
To show left to right suppose there is a Λ-consistent set Γ. Then $\Gamma \nvdash_\Lambda \bot$ by definition. So, because Λ is strongly complete w.r.t. $\mathbb{F}$, $\Gamma \nvDash_{\mathbb{F}} \bot$. By definition this means that there is a $\mathcal{F} \in \mathbb{F}$ and a $w \in \mathcal{F}$ and a valuation $\mathcal{V}$ such that $(\mathcal{F}, \mathcal{V}), w \models \Gamma$ and $(\mathcal{F}, \mathcal{V}), w \nvDash \bot$. Especially $(\mathcal{F}, \mathcal{V}), w \models \Gamma$ is proven. □

Corollary 1.2.10 *Λ is weakly complete w.r.t. a class of structures $\mathbb{F}$ iff every Λ-consistent formula is satisfiable on some $F \in \mathbb{F}$.*

Proof. Use 1.2.6 and 1.2.9. □

Soundness was, as we have seen in the proof of 1.2.4, rather easy to show. Completeness on the other hand is more difficult. By 1.2.9 one has to build (suitable) models such that every Λ-consistent set of formulas can be satisfied in these models.
One good technique is to build models out of maximal consistent sets of formulas. This approch leads to *canonical models* and is probably the most basic algorithm to show (strong) completeness. However, it does not work all the time, but it suffices for our purposes. See [1] for more details on this.
Canonical models show strong completeness, so once we have established the completeness results for the modal logics defined in 1.2.3 we can use 1.2.6 to obtain a weak completeness result. With this, we will have linked those logics

to a class of frames (see [10]) and therefore have completed our task.

1.3 Canonical Models

As said before, we want to build models such that for some given logic Λ every Λ-consistent set of formulas can be satisfied in these models. And building these models from maximal consistent sets is a very good way to do so, as will become apparent in this section. But first we have to define what is meant by a maximal consistent set.

Definition 1.3.1 A set of formulas Γ is maximal Λ-consistent for a logic Λ if Γ is Λ-consistent and any Θ with $\Gamma \subsetneq \Theta$ is Λ-inconsistent.
Maximal Λ-consistent sets are called Λ-MCS's or just MCS's when Λ is unambiguous from the context.

Such MCS's have some nice properties.

Lemma 1.3.2 *Let Λ be a logic and Γ a Λ-MCS.*

(i) for all formulas ϕ: $\phi \in \Gamma$ or $\neg\phi \in \Gamma$

(ii) Γ is closed under modus ponens

(iii) $\Lambda \subseteq \Gamma$

(iv) for all formulas ϕ, ψ: $\phi \vee \psi \in \Gamma$ iff $\phi \in \Gamma$ or $\psi \in \Gamma$

Proof. Let Γ, Λ, ϕ and ψ be as above. Then

(i) It can not be that both $\Gamma \cup \{\phi\}$ and $\Gamma \cup \{\neg\phi\}$ are inconsistent, because if it were so we would have $\Gamma \vdash_\Lambda \phi$ and $\Gamma \vdash_\Lambda \neg\phi$ by the contradiction lemma 1.2.8. But this would mean that Γ is inconsistent, a contradiction! So let $\Gamma \cup \{\phi\}$ be consistent. By maximality of Γ: $\Gamma = \Gamma \cup \{\phi\}$, so $\phi \in \Gamma$.

(ii) Let $\phi \in \Gamma$ and $(\phi \to \psi) \in \Gamma$ and suppose $\psi \notin \Gamma$. So by (i) $\neg\psi \in \Gamma$. But $\phi \wedge (\phi \to \psi) \wedge \neg\psi \vdash_\Lambda \bot$, in contradiction to the consistency of Γ!

(iii) Supposing the opposite and using (i) gives a $\phi \in \Lambda$ with $\neg\phi \in \Gamma$, which contradicts consistency.

(iv) Suppose $\phi \vee \psi \in \Gamma$ and $\phi \notin \Gamma$ and $\psi \notin \Gamma$. By (i) the negations are in Γ. But $(\phi \vee \psi) \wedge \neg\phi \wedge \neg\psi \vdash_\Lambda \bot$, a contradiction to consistency.
For right to left suppose $\phi \in \Gamma$ and ψ arbitrary and $\phi \vee \psi \notin \Gamma$. Then $\neg(\phi \vee \psi) \in \Gamma$, i.e. $(\neg\phi \wedge \neg\psi) \in \Gamma$ which leads to inconsistency of Γ by $\phi \in \Gamma$.

[10] The class is not unique. For example [5] shows that S4.2 is also weakly complete w.r.t. the class of baled trees and the class of finite pre-lattices.

□

Definition 1.3.3 The canonical model $\mathcal{M}^\Lambda$ for a normal modal logic Λ is a triple $(\mathcal{W}^\Lambda, \mathcal{R}^\Lambda, \mathcal{V}^\Lambda)$ where

(i) $\mathcal{W}^\Lambda$ is the set of all Λ-MCSs.

(ii) $\mathcal{R}^\Lambda$ is a binary relation on $\mathcal{W}^\Lambda$ defined by $(w,u) \in \mathcal{R}^\Lambda$ if for all formulas ϕ, $\phi \in u$ implies $\Diamond\phi \in w$.
$\mathcal{R}^\Lambda$ is called the canonical relation.

(iii) $\mathcal{V}^\Lambda$ is the valuation defined by $\mathcal{V}^\Lambda(p) = \{w \in \mathcal{W}^\Lambda \mid p \in w\}$.
$\mathcal{V}^\Lambda$ is called the canonical valuation.

$\mathcal{F}^\Lambda = (\mathcal{W}^\Lambda, \mathcal{R}^\Lambda)$ is called the canonical frame of Λ.

As we will see in the Truth Lemma 1.3.6, a formula is true in the canonical model at w iff the formula is an element of w. So to satisfy every consistent set of formulas, it would be very convenient to just enrich the set to an MCS. That this is in fact possible is the statement of Lindenbaum's Lemma:

Lemma 1.3.4 (Lindenbaum's Lemma) *Let Λ be a logic. If Σ is a Λ-consistent set of formulas, then there is a Λ-MCS Σ^+ with $\Sigma \subseteq \Sigma^+$.*

Proof. Let $\phi_0, \phi_1, ...$ be an enumeration of the modal formulas. Let

$$\Sigma_0 = \Sigma$$
$$\Sigma_{n+1} = \begin{cases} \Sigma_n \cup \{\phi\} & \text{if this is } \Lambda \text{ -consistent} \\ \Sigma_n \cup \{\neg\phi\} & \text{otherwise} \end{cases}$$
$$\Sigma^+ = \bigcup_{n \geq 0} \Sigma_n$$

Claim 1. Σ_n is Λ-consistent for all n
Proof of Claim 1. Ind(n)
For $n = 0$ this is true by the hypothesis that Σ is Λ-consistent. For $n + 1$, Σ_{n+1} is either $\Sigma_n \cup \{\phi_n\}$ and Λ-consistent by definition, or $\Sigma_{n+1} = \Sigma_n \cup \{\neg\phi_n\}$. In this case $\Sigma_n \cup \{\phi_n\}$ is inconsistent and by the contradiction lemma 1.2.8: $\Sigma_n \vdash_\Lambda \neg\phi_n$. By induction hypothesis Σ_n is consistent, hence $\Sigma_n \cup \{\neg\phi_n\}$ is consistent. □(Claim 1)

Claim 2. For every formula ϕ either $\phi \in \Sigma^+$ or $\neg\phi \in \Sigma^+$.
Proof of Claim 2. $\phi \in \Sigma^+$ and $\neg\phi \in \Sigma^+$ implies $\phi \in \Sigma_n$ and $\neg\phi \in \Sigma_n$ for some n, and hence inconsistency of Σ_n, a contradiction to Claim 1.
Furthermore we find an m such that $\phi = \phi_m$. So $\phi \in \Sigma_m$ or $\neg\phi \in \Sigma_m$. The claim follows. □(Claim 2)

Claim 3. Σ^+ is a Λ-MCS.
Proof of Claim 3. If Σ^+ were inconsistent we would find a finite (!) sequence $\phi^1, ..., \phi^n$ such that $\vdash_\Lambda \phi^1 \wedge ... \wedge \phi^n \rightarrow\perp$. But such a sequence would be in Σ_k for some $k < \omega$ big enough, contradicting Claim 1.
For maximality let $\Sigma^+ \subsetneq \Gamma$ for some set of formulas Γ and let ϕ be a witness. By Claim 2 $\neg\phi \in \Sigma^+$ so by definition $\Sigma^+ \vdash_\Lambda \neg\phi$. By $\Sigma^+ \subsetneq \Gamma$: $\Gamma \vdash_\Lambda \neg\phi$. But $\phi \in \Gamma$ by hypothesis, hence Γ is inconsistent. □(Claim 3)

□

The next lemma states that there are enough MCS for our purposes.

Lemma 1.3.5 (Existence Lemma) *For any normal modal logic Λ and any state $w \in \mathcal{W}^\Lambda$, if $\Diamond\phi \in w$ then there is a state $v \in \mathcal{W}^\Lambda$ such that $w\mathcal{R}^\Lambda v$ and $\phi \in v$.*

Proof. Let $\Diamond\phi \in w$ and define $v^- = \{\phi\} \cup \{\psi \mid \Box\psi \in w\}$.
Claim 1. v^- is consistent
Proof of Claim 1. Suppose not. Then by the Contradiction Lemma 1.2.8 $\{\psi \mid \Box\psi \in w\} \vdash_\Lambda \neg\phi$, i.e. there are $\psi_1, ..., \psi_n$ such that $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \rightarrow \neg\phi$. So by (nec) $\vdash_\Lambda \Box((\psi_1 \wedge ... \wedge \psi_n) \rightarrow \neg\phi)$. By (K) $\vdash_\Lambda \Box(\psi_1 \wedge ... \wedge \psi_n) \rightarrow \Box\neg\phi$. Now $\Box(\psi_1 \wedge ... \wedge \psi_n) \rightarrow (\Box\psi_1 \wedge ... \wedge \Box\psi_n)$ is a theorem of every normal modal logic, so $\vdash_\Lambda (\Box\psi_1 \wedge ... \wedge \Box\psi_n) \rightarrow \Box\neg\phi$. Since w is a MCS and $\Box\psi_i \in w$ for $i < n$, the conjunction is also in w. By consistency $\Box\neg\phi$ is also in w and hence by (Dual) $\neg\Diamond\phi \in w$, a contradiction to $\Diamond\phi \in w$ and w is MCS. □(Claim 1)

Now let v be an MCS-extension of v^-, whose existence is granted by Lindenbaum's Lemma. Then $\phi \in v$ by construction and for all ψ: $\Box\psi \in w$ implies $\psi \in v$. This yields $w\mathcal{R}^\Lambda v$ by the following argument:
Let $\phi \in v$ and suppose $\Diamond\phi \notin w$. Because w is a MCS $\neg\Diamond\phi \in w$. Hence $\neg\Diamond\neg\neg\phi \in w$ and by (Dual) $\Box\neg\phi \in w$. It follows that $\neg\phi \in v$, a contradiction to v being a MCS (see [11]). □

We now have everything we need to establish the "truth = membership" equation.

Lemma 1.3.6 (Truth Lemma) *For any normal modal logic Λ and any formula ϕ*

$$\mathcal{M}^\Lambda, w \models \phi \textit{ iff } \phi \in w.$$

Proof. By induction on the complexity of ϕ.
The base case follows from definition. The Boolean Cases follow from 1.3.2.

[11] The reverse of the last argument is also true: Let $w\mathcal{R}^\Lambda v$ and $\phi \notin v$. Then $\neg\phi \in v$ and therefore $\Diamond\neg\phi \in w$. w consistently yields $\neg\Diamond\neg\phi \notin w$, i.e., $\Box\phi \notin w$. So $\Box\phi \in w \Rightarrow \phi \in v$ holds for all ϕ.

We have $\mathcal{M}^\Lambda, w \models \Diamond\phi$ iff
$\exists v(w\mathcal{R}^\Lambda v \wedge \mathcal{M}^\Lambda, v \models \phi)$ iff (by induction hypothesis)
$\exists v(w\mathcal{R}^\Lambda v \wedge \phi \in v)$ implies (by def of $\mathcal{R}^\Lambda$) $\Diamond\phi \in w$.
For the reverse direction, suppose $\Diamond\phi \in w$. By the above equivalence it suffices to find some v such that v is a MCS, $w\mathcal{R}^\Lambda v$ and $\phi \in v$. But this is precisely what the Existence Lemma gives us. □

We have the definitions and properties needed to work out some completeness results. The next theorem puts together what we have constructed and analysed so far in this section.

Theorem 1.3.7 (Canonical Model Theorem) *Any normal modal logic is strongly complete with respect to its canonical model.*

Proof. Given a consistent set of formulas Σ of the modal logic Λ. By Lindenbaum's Lemma we find some MCS Σ^+ such that $\Sigma \subseteq \Sigma^+$. Now by the Truth Lemma above: $\mathcal{M}^\Lambda, \Sigma^+ \models \Sigma$. The theorem follows by 1.2.9. □

At first glance, the Canonical Model Theorem 1.3.7 seems rather abstract and not very helpful. To show completeness for some logic Λ w.r.t. the singleton class of its canonical model does not seem useful. But by lemma 1.2.9 we only have to show that the canonical model is an element of the class considered. Then for every consistent set Γ of this logic, the canonical model contains a state Γ^+ extending Γ such that Γ holds in the canonical model at state Γ^+. By 1.2.9 we have shown strong completeness.

To put this to work, the following lemma shows slightly more than we actually need.

Lemma 1.3.8 *For any normal modal logic Λ with canoncal model $\mathcal{M}^\Lambda$*

(i) (4) $\in \Lambda \Rightarrow \mathcal{R}^\Lambda$ is transitive.

(ii) (S) $\in \Lambda \Rightarrow \mathcal{R}^\Lambda$ is reflexive.

(iii) (S), (4),(.2) $\in \Lambda \Rightarrow \mathcal{R}^\Lambda$ is a directed partial order.

(iv) (5) $\in \Lambda \Rightarrow \mathcal{R}^\Lambda$ is symmetrical.

Proof. (i) Suppose $w, v, u \in \mathcal{M}^\Lambda$ such that $w\mathcal{R}^\Lambda v$ and $v\mathcal{R}^\Lambda u$. We need to show that $w\mathcal{R}^\Lambda u$. Suppose $\phi \in u$. By $v\mathcal{R}^\Lambda u$ we have $\Diamond\phi \in v$ and by $w\mathcal{R}^\Lambda v$ $\Diamond\Diamond\phi \in w$. But (4)$\in \Lambda$ and w is a Λ-MCS. So by modus ponens $\Diamond\phi \in w$ and hence $w\mathcal{R}^\Lambda u$.

(ii) By (S)$\in \Lambda$ we have $(\phi \to \Diamond\phi) \in \Lambda$ (see [12]). Let w be a world in $\mathcal{M}^\Lambda$ with

[12] *Proof.* $\Box p \to p$ iff (Dual) $\neg(\neg\Diamond\neg p) \vee p$ iff $\Diamond\neg p \vee p$ iff (uni.sub.) $\Diamond q \vee \neg q$ iff $q \to \Diamond q$ □

$\phi \in w$. Because w is a Λ-MCS and $(\phi \to \Diamond\phi) \in \Lambda$ it follows from modus ponens that $\Diamond\phi \in w$. Thus $w\mathcal{R}^\Lambda w$.

(iii) It only remains to show directedness, i.e.

$$\forall x, y, z(x\mathcal{R}^\Lambda y \wedge x\mathcal{R}^\Lambda z \to \exists u(y\mathcal{R}^\Lambda u \wedge z\mathcal{R}^\Lambda u)).$$

Suppose $w_0, w_1, w_2 \in \mathcal{W}^\Lambda$ with $w_0\mathcal{R}^\Lambda w_1$ and $w_0\mathcal{R}^\Lambda w_2$. Define

$$w_3 = \{\phi \mid \Box\phi \in w_1\} \cup \{\neg\psi \mid \Diamond\psi \notin w_2\}.$$

We will see that w_3 is Λ-consistent which enables us to use Lindenbaum's Lemma to enrich the set to a maximal Λ-consistent set. This enriched set will be $\mathcal{R}^\Lambda$-greater than w_1, w_2.

Suppose w_3 is not Λ-consistent. So there are ϕ_i and ψ_j with $i \leq n$ and $j \leq m$ such that $\Box\phi_i \in w_1$, $\Diamond\psi_j \notin w_2$ and $\vdash_\Lambda (\bigwedge_{i\leq n} \phi_i \wedge \bigwedge_{j\leq m} \neg\psi_j) \to\bot$. Set $\phi = \bigwedge_{i\leq n} \phi_i$ and $\neg\psi = \bigwedge_{j\leq m} \neg\psi_j$. Then $\phi \wedge \neg\psi$ implies $\bot$ in Λ, so $\phi \to \psi$ is an element of every maximal Λ consistent set. So $\phi \to \psi$ is in w_1. By distribution of "$\Box$" over "$\wedge$": $\Box\phi \in w_1$. This implies $\Box\psi \in w_1$ because if $\Box\psi \notin w_1$ then $\neg\Box\psi = \Diamond\neg\psi \in w_1$, i.e. $\exists v\ w_1\mathcal{R}^\Lambda v$ and $\phi \wedge \neg\psi \in v$, a contradiction to the consistency of v. Hence $\Box\psi \in w_1$, i.e. $\Diamond\Box\psi \in w_0$. By (.2) also $\Box\Diamond\psi \in w_0$ which implies $\Diamond\psi \in w_2$, a contradiction!

Let w_3^+ be a maximal consistent enrichment of w_3. By $\Box\phi \in w_1 \Rightarrow \phi \in w_3^+$ we have $w_1\mathcal{R}^\Lambda w_3^+$. Also $\Diamond\psi \notin w_2$ implies $\psi \notin w_3^+$, i.e. $\psi \in w_3^+ \Rightarrow \Diamond\psi \in w_2$, so $w_2\mathcal{R}^\Lambda w_3^+$.

(iv) By (5)$\in \Lambda$: $(\phi \to \Box\Diamond\phi) \in \Lambda$ (see [13]). Suppose w is a world in $\mathcal{M}^\Lambda$ and $\phi \in w$. By the above argument $\Box\Diamond\phi \in w$. Now suppose a world v in $\mathcal{M}^\Lambda$ such that $w\mathcal{R}^\Lambda v$. By the footnote in the proof of the Existence Lemma 1.3.5 $\Box\Diamond\phi \in w \wedge w\mathcal{R}^\Lambda v$ implies $\Diamond\phi \in v$. By $\phi \in w$ it follows that $v\mathcal{R}^\Lambda w$ as desired.

□

It is now just a matter of selecting and combining the right pieces to establish the completeness results we will need during the further discussion of the modal logic of forcing.

Theorem 1.3.9 *We have the following completeness results:*

(i) K is strongly complete with respect to the class of all frames.

(ii) S4 is strongly complete with respect to the class of all reflexive and transitive frames.

[13]

Proof. $\Diamond\Box p \to p$ iff (Dual) $\neg(\neg\neg\Diamond\neg p) \vee p$ iff $\Box\Diamond\neg p \vee p$ iff (uni.sub.) $\Box\Diamond q \vee \neg q$ iff $q \to \Box\Diamond q$) □

(iii) S4.2 is strongly complete with respect to the class of all frames whose order is a dirceted partial order.

(iv) S5 is strongly complete with respect to the class of frames whose relation is an equivalence relation.

Proof. Given a set of formulas Γ consistent for the logic in question. By 1.2.9 it then will do to find a model $(\mathcal{F},\mathcal{V})$ such that:

a) there is a state w in $(\mathcal{F},\mathcal{V})$ such that $(\mathcal{F},\mathcal{V}), w \models \Gamma$.

b) $\mathcal{F}$ satisfies the claimed attribute(s).

Now the canonical model satisfies, at the state Γ^+, where Γ^+ is a MCS extending Γ, condition a). So it remains to show that b) holds for the canonical frame. This was already done by the preceding lemma. $\square$

Putting 1.3.9 and 1.2.4 together, we have shown that the considered logics are sound and (strongly) complete w.r.t. the given class of frames.

We now introduce a notion which will help us to see that we can demand an additional property of the frames, namely that the frames are *rooted*.

Definition 1.3.10 Let $\mathcal{M}$=($\mathcal{W}$, $\mathcal{R}$, $\mathcal{V}$) be a model. A model $\mathcal{N}$=($\mathcal{W}'$, $\mathcal{R}'$, $\mathcal{V}'$) is called a *submodel* of $\mathcal{M}$, $\mathcal{N} \subseteq \mathcal{M}$, if

- $\mathcal{W}' \subseteq \mathcal{W}$
- $\mathcal{R}' = \mathcal{R} \upharpoonright \mathcal{W}'$
- $\forall p \in \Phi : \mathcal{V}'(p) = \mathcal{V}(p) \cap \mathcal{W}'$

$\mathcal{N}$ is called a *generated submodel of* $\mathcal{M}$, $\mathcal{N} \sqsubseteq \mathcal{M}$, if $\mathcal{N} \subseteq \mathcal{M}$ and $\mathcal{W}'$ is an upwards closed subset of $\mathcal{W}$.
$\mathcal{N}$ is called *generated by the set* X if

$$\mathcal{W}' = \{x \in \mathcal{W} \mid \exists y \in X : y\mathcal{R}x\}$$

and *point generated* if X is a singleton.

The following theorem is easy to proof but still very helpful.

Theorem 1.3.11 (Generation Theorem) *Suppose $\mathcal{N} = (\mathcal{W}', \mathcal{R}', \mathcal{V}')$ is a generated submodel of $\mathcal{M} = (\mathcal{W}, \mathcal{R}, \mathcal{V})$. Then for every formula ϕ and every $x \in \mathcal{W}'$*

$$\mathcal{N}, x \models \phi \text{ iff } \mathcal{M}, x \models \phi$$

Proof. Ind(ϕ)
The atomic case follows from definition.
The case $\phi = \phi_1 \wedge \phi_2$ is also trivial.
Let $\phi = \neg\psi$. Then $\mathcal{N}, x \models \phi$ iff $\mathcal{N}, x \not\models \psi$ iff (by induction hypothesis) $\mathcal{M}, x \not\models \psi$

iff $\mathcal{M}, x \models \phi$.
Now let $\phi = \Diamond\psi$.
$\mathcal{N}, x \models \phi$ iff
$\exists y \in \mathcal{W}' : x\mathcal{R}'y \wedge \mathcal{N}, y \models \psi$ iff (by induction hypothesis, note that we use that $\mathcal{W}'$ is upwards closed in $\mathcal{W}$ and $x \in \mathcal{W}'$ for the direction from bottom to top)
$\exists y \in \mathcal{W} : x\mathcal{R}y \wedge \mathcal{M}, y \models \psi$ iff
$\mathcal{M}, x \models \phi$. □

The Generation Theorem yields that a normal modal logic Λ which is sound and complete w.r.t. some class of frames satisfying some property P is also sound and complete w.r.t. the class of rooted frames with property P. We will show this for S4.2 in the next lemma.

Definition 1.3.12 A frame is called *rooted* if it has a unique minimal element.

Corollary 1.3.13 *Every point generated submodel is rooted.*

Lemma 1.3.14 *S4.2 is sound and complete w.r.t. the class of rooted frames whose order is a directed partial order.*

Proof. Let $\mathbb{F}$ be the class of frames whose order is a directed partial order and $\mathbb{F}_r \subset \mathbb{F}$ be the subclass of rooted frames in $\mathbb{F}$. We need to see that $S4.2 = \Lambda_{\mathbb{F}_r}$. Suppose $\phi \in S4.2$ and assume that there is some $\mathcal{F}_r \in \mathbb{F}_r$ such that $\mathcal{F}_r \not\models \phi$. By $\mathbb{F}_r \subset \mathbb{F}$ we have $\mathcal{F}_r \in \mathbb{F}$, contradicting the soundness result from lemma 1.2.4. On the other hand: $\phi \notin S4.2$ and $S4.2 = \Lambda_{\mathbb{F}}$ gives us some $\mathcal{F} \in \mathbb{F}$ and some $w \in \mathcal{F}$ such that for any model $\mathcal{M}$ over $\mathcal{F}$: $\mathcal{M}, w \not\models \phi$. Fix such an $\mathcal{M}$. Now let $\mathcal{N}$ be the submodel of $\mathcal{M}$ that was generated by w. Then $\mathcal{N}, w \not\models \phi$ by the Generation Theorem. Hence $\phi \notin \Lambda_{\mathbb{F}_r}$. □

1.4 Filtration and the Finite Frame Property

The canonical model for a given (consistent and normal) logic Λ refutes every ϕ with $\nvdash_\Lambda \phi$ at some world. While this property is very nice, the model itself is very big and clumsy (it contains continuum many worlds). We would like to have a smaller model to handle it more elegantly. This can, for some logics, in fact be done, as this section will show.

Definition 1.4.1 A logic Λ has the *finite frame property* if for every $\phi \notin \Lambda$ there is a finite frame $\mathcal{F}$ such that $\mathcal{F}, w \not\models \phi$ for some world w in $\mathcal{F}$ and $\mathcal{F} \models \Lambda$.

We will see that S4.2 has the finite frame property by using a method that is called filtration. The idea is to take any model that refutes every ϕ not in S4.2 at some world (those models exist, the canonical model is an example) and a (well-behaved) set of formulas. Then we define an equivalence class on this

model by using the set. When we take the "quotient" of the model through the equivalence class, the resulting model will be called the filtration of the model through the set (assuming some properties for our equivalence class). The Filtration Theorem then states that every formula in the given set holds in the model if and only if it holds in the filtration, i.e. if ϕ is in the set of formuals, then the filtration falsifies it at some world. Because the filtration will be finite and its order is a directed partial order, we have shown that S4.2 has the finite frame property.
This section uses the notation from [2] and follows their construction very closely.

Let $\mathcal{M} = (\mathcal{F}, \mathcal{V})$ be a model, $\mathcal{F} = (\mathcal{W}, \mathcal{R})$, and Σ be a set of formulas which is closed under subformulas. Given two worlds $x, y \in \mathcal{W}$ we say that x is Σ-equivalent to y in $\mathcal{M}$, $x \sim_\Sigma y$, if

$$\mathcal{M}, x \models \phi \text{ iff } \mathcal{M}, y \models \phi \text{ for every } \phi \in \Sigma$$

Note that $\sim_\Sigma$ is trivially an equivalence relation on $\mathcal{W}$, so we denote the equivalence class of x by $[x]_\Sigma = \{y \in \mathcal{W} \mid x \sim_\Sigma y\}$. We will drop the subscribt Σ if this does not involve ambiguity.

Definition 1.4.2 A *filtration of $\mathcal{M}$ through* Σ is any model $\mathcal{N} = (\mathcal{G}, \mathcal{U})$ based on a frame $\mathcal{G} = (\mathcal{O}, \mathcal{S})$ such that

(i) $\mathcal{O} = \{[x] \mid x \in \mathcal{W}\}$

(ii) $\mathcal{U}(p) = \{[x] \mid x \in \mathcal{V}(p)\}$ for every variable $p \in \Sigma$

(iii) $x\mathcal{R}y \Rightarrow [x]\mathcal{S}[y]$ for all $x, y \in \mathcal{W}$

(iv) if $[x]\mathcal{S}[y]$ then $\mathcal{M}, x \models \Box\phi$ implies $\mathcal{M}, y \models \phi$ for $x, y \in \mathcal{W}$ and $\Box\phi \in \Sigma$

Remark 1.4.3 In general the conditions (*iii*) and (*iv*) do not determine $\mathcal{S}$ uniquely. In fact one can show that $\underline{\mathcal{S}} \subseteq \mathcal{S} \subseteq \overline{\mathcal{S}}$ for any filtration where

$$\overline{\mathcal{S}} = \{([x], [y]) \mid \forall \Box\phi \in \Sigma(\mathcal{M}, x \models \phi \to \mathcal{M}, y \models \phi)\}$$
$$\underline{\mathcal{S}} = \{([x], [y]) \mid \exists x', y' \in \mathcal{W}(x \sim x' \wedge y \sim y' \wedge x'\mathcal{R}y'\}$$

$\underline{\mathcal{S}}$ is called the finest filtration of $\mathcal{M}$ through Σ, while $\overline{\mathcal{S}}$ is called the coarsest. We will only use the finest filtration, so we do not bother with a proof of the above and its implications. The interested reader is referred to [2], proposition 5.27 and page 141 and following.

Corollary 1.4.4 *With the notation as above:* $|\mathcal{G}| \leq 2^{|\Sigma|}$.

Proof. Define

$$\pi : \mathcal{G} \to {}^{\Sigma}2$$
$$[x] \mapsto f_x$$

where

$$f_x(\phi) = \begin{cases} 1 & \text{if } \mathcal{M}, x \models \phi \\ 0 & \text{otherwise} \end{cases}$$

To see that π is well-defined let $y, z \in [x]$ where $[x] \in \mathcal{G}$. For all $\phi \in \Sigma$ we use the definition of the equivalence relation to obtain: $f_y(\phi) = 1$ iff $\mathcal{M}, y \models \phi$ iff $\mathcal{M}, x \models \phi$ iff $\mathcal{M}, z \models \phi$ iff $f_z(\phi) = 1$.

This shows that π is not dependent on the choice of the member of an equivalence class and is hence well-defined.

Now let $f_x \neq f_y$. Without loss of generality there then is some $\phi \in \Sigma$ where $f_x(\phi) = 1$ and $f_y(\phi) = 0$. So $\mathcal{M}, x \models \phi$ and $\mathcal{M}, y \models \neg\phi$. Because $\phi \in \Sigma$ this means $[x] \neq [y]$, i.e. π is injective. This proves the corollary.

□

We now show that a filtration is well-behaved.

Theorem 1.4.5 (Filtration Theorem) *Let $\mathcal{M}$ be a model and $\mathcal{N}$ a filtration of $\mathcal{M}$ through a set of formulas Σ that is closed under subformulas. Then for every world $x \in \mathcal{M}$ and every formula $\phi \in \Sigma$*

$$\mathcal{M}, x \models \phi \textit{ iff } \mathcal{N}, [x] \models \phi$$

Proof. By induction on the complextity of ϕ.
The base case follows from (ii) from the above definition.
The Boolean combinations follow directly from the truth-definitions and the induction hypothesis.
Let $\Box\phi \in \Sigma$ and $\mathcal{M}, x \models \Box\phi$. We need to show that $\mathcal{N}, [y] \models \phi$ for every $[y]$ with $[x]\mathcal{S}[y]$. Take such a $[y]$. By (iv) we see $\mathcal{M}, y \models \phi$, so by induction hypothesis $\mathcal{N}, [y] \models \phi$ (note that $\phi \in \Sigma$ by the closure properties of Σ). Conversely suppose $\mathcal{N}, [x] \models \Box\phi$. Let y be such that $x\mathcal{R}y$. So $[x]\mathcal{S}[y]$ by (iii). Therefore $\mathcal{N}, [y] \models \phi$ and hence, by induction hypothesis and closure properties, $\mathcal{M}, y \models \phi$. By arbitrarity of y: $\mathcal{M}, x \models \Box\phi$. □

Lemma 1.4.6 *S4.2 has the finite frame property.*

Proof. Suppose $\phi \notin$S4.2. By 1.3.14 there is a model $\mathcal{M} = (\mathcal{W}, \mathcal{R}, \mathcal{V})$ whose relation is a rooted directed partial order such that there is a world w with $\mathcal{M}, w \not\models \phi$. Let $\Gamma = \Gamma_\phi = \{\psi \mid \psi \text{ is subformula of } \phi\}$.
Let $\mathcal{N}$ be the finest filtration of $\mathcal{M}$ through Γ and let $\mathcal{N}^t$ be its transitive closure.
Claim 1. $\mathcal{N}^t$ is a filtration of $\mathcal{M}$ through Γ.
Proof of Claim 1. Suppose that $\mathcal{S}$ is the relation on $\mathcal{N}$. The relation on $\mathcal{N}^t$ is

$$\mathcal{S}^t = \{([x], [y]) \mid \exists n > 0 \ [x]\mathcal{S}^n[y]\}$$

where $x\mathcal{S}^n y$ iff $\exists x_0, ..., x_{n-1} \ x = x_0 \wedge x_0\mathcal{S}x_1 \wedge ... \wedge x_{n-1}\mathcal{S}y$.
Note that condition (i) and (ii) were not affected by making $\mathcal{S}$ transitive, so they hold in $\mathcal{N}^t$.

For (iii) we have $x\mathcal{R}y \implies [x]\mathcal{S}[y] \implies [x]\mathcal{S}^t[y]$, so the condition is satisfied.
To see (iv) suppose $[x]\mathcal{S}^t[y]$ and $\mathcal{M}, x \models \Box\psi$ for some $x, y \in \mathcal{W}$ and $\Box\psi \in \Gamma$. Then there is a finite sequence $[u], ..., [v]$ such that $[x]\mathcal{S}[u]\mathcal{S}...\mathcal{S}[v]\mathcal{S}[y]$. Because $\mathcal{S}$ is the finest filtration there are x', u' such that $x' \sim x$ and $u' \sim u$ and $x'\mathcal{R}u'$. Because $\mathcal{R}$ is transitive we know $\mathcal{M}, u' \models \Box\psi$. Since $\Box\psi$ was in Γ and $u \sim u'$ we have $\mathcal{M}, u \models \Box\psi$. Using this argument finitely often we eventually see $\mathcal{M}, v \models \Box\psi$, so, by (iv) for $\mathcal{S}$ and $[v]\mathcal{S}[y]$, we see $\mathcal{M}, y \models \psi$ as desired. $\Box$(Claim 1)

By the Filtration Theorem we now have

$$\mathcal{M}, x \models \xi \text{ iff } \mathcal{N}^t, [x] \models \xi \text{ for all } \xi \in \Gamma,\ x \in \mathcal{W}.$$

Therefore $\mathcal{N}^t, w \not\models \phi$.
Since every formula is finite, the set Γ_ϕ is finite and hence, by 1.4.4, $\mathcal{N}^t$ is finite. It only remains to show that $\mathcal{N}^t$ has a S4.2 frame. Because S4.2 is sound and complete with respect to the class of directed partial orders it is sufficient to show that $\mathcal{S}^t$ is such a order. Now, $\mathcal{S}^t$ is transitive by definition. It is also reflexive because $\mathcal{R}$ was reflexive, i.e. $\forall x \in \mathcal{W} : x\mathcal{R}x$ which implies $[x]\mathcal{S}^t[x]$ for all $[x]$ by (iii).
To see that $\mathcal{S}^t$ is directed let $[x]\mathcal{S}^t[y]$ and $[x]\mathcal{S}^t[z]$ for some worlds $[x], [y], [z]$ in $\mathcal{N}^t$. If $[y] = [z]$ we are done by reflexivity, so suppose $[y] \neq [z]$. Then there are $u, v \in \mathcal{W}$ and $y', z' \in \mathcal{W}$ such that $y \sim y'$, $z \sim z'$ and $u\mathcal{R}y'$ and $v\mathcal{R}z'$ because $\mathcal{S}^t$ includes the finest filtration. Clearly $y' \neq z'$. Because $\mathcal{M}$ had a rooted and transitive frame we have for the root r that $r\mathcal{R}y' \wedge r\mathcal{R}z'$ and since $\mathcal{R}$ is directed there is a w such that $y'\mathcal{R}w \wedge z'\mathcal{R}w$ which implies $[y'] = [y]\mathcal{S}^t[w]$ and $[z'] = [z]\mathcal{S}^t[w]$ by (iii). Hence $\mathcal{S}^t$ is directed. $\Box$

The following lemma reformulates the characterisation of S4.2 and provides us with the characterisation we will use throughout the rest of the thesis. Note that a lattice is a partially ordered set such that every two elements of the lattice have a unique least upper bound and a unique greatest lower bound, called join and meet respectively. A pre-lattice is a partially ordered set $\mathcal{F}$ such that $\mathcal{F}/\equiv$ forms a lattice, where $\forall a, b \in \mathcal{F} : a \equiv b$ iff $a \leq_{\mathcal{F}} b \leq_{\mathcal{F}} a$.

Lemma 1.4.7 *S4.2 is sound and complete with respect to the class of frames whose relation is a finite pre-lattice.*

In order to prove this, we will take a model with a directed partial order frame and construct a new model out of it whose relation is a finite pre-lattice and which is "similar" to the model we started with. The construction is essentially "tree-unravelling", a standard method used to obtain trees from partial orders, see the proof of theorem 2.19 in [2]. But first we state what is meant by "similar":

Definition 1.4.8 Two Kripke-models $\mathcal{M} = (\mathcal{W}, \mathcal{R}, \mathcal{V})$ and $\mathcal{M}' = (\mathcal{W}', \mathcal{R}', \mathcal{V}')$ are called *bisimilar* if there is a correspondence of their worlds $a \sim a'$ for $a \in \mathcal{M}, a' \in \mathcal{M}'$ such that if $a \sim a'$ then

- $a \in \mathcal{V}(p)$ iff $a' \in \mathcal{V}'(p)$ for every propositional variable p
- $a\mathcal{R}b \Rightarrow \exists b' \in \mathcal{M}'\ a'\mathcal{R}'b'$ and $b' \sim b$
- $a'\mathcal{R}b' \Rightarrow \exists b \in \mathcal{M}\ a\mathcal{R}b$ and $b \sim b'$

Note that it follows from induction that corresponding worlds have exactly the same modal truths. Also note that the correspondence is not required to be functional or one-to-one.

With this definition at hand we can now start the missing proof.

Proof of lemma. Let $\phi \notin S4.2$. Because S4.2 has the finite frame property by 1.4.6 there is a model $\mathcal{M} = (\mathcal{F}, \mathcal{V})$ such that ϕ fails at some world w_0 in $\mathcal{M}$ and the relation in $\mathcal{M}$ is a directed partial order. Without loss of generality we may assume that w_0 is in the smallest cluster (see [14]) of $\mathcal{F}$. By directedness and because $\mathcal{F}$ is finite there is a largest cluster in $\mathcal{F}$. Call it $[z]$.

Recall that $\mathcal{F}/\equiv$ is a finite directed partial order where $a \equiv b$ iff $a \leq_{\mathcal{F}} b \leq_{\mathcal{F}} a$. Therefore, for any $[x] \in \mathcal{F}/\equiv$, the interval $[[w_0], [x]]$ is a union of linearly ordered sets. We define

t is a path from $[w_0]$ to $[x]$ if t is a maximal linearly ordered subset of $[[w_0], [x]]$.

Then

$$T = \{t \mid \exists x \text{ such that } t \text{ is a path from } [w_0] \text{ to } [x]\}$$

is a tree if ordered by end-extension.

We now define the bisimilar model as follows:

$\mathcal{W}' = \{(x,t) \mid t \text{ is path from } [w_0] \text{ to } [x] \text{ in } \mathcal{F}/\equiv \text{ and } [x] \neq [z]\} \cup \{[z]\}$

We order the elements by end-extension of the paths and the order in $\mathcal{F}/\equiv$ with $[z]$ still maximal, i.e.

$\mathcal{R}' = \{((x,t),(y,s)) \mid$
$(x,t),(y,s) \in \mathcal{W}' \wedge t \subseteq s \wedge x \leq_{\mathcal{F}/\equiv} y\} \cup \{((x,t),[z]) \mid (x,t) \in \mathcal{W}'\}$

Now define a valuation on $\mathcal{F}' = (\mathcal{W}', \mathcal{R}')$ by

$\mathcal{V}'(p) = \{(x,t) \mid x \in \mathcal{V}(p)\}$

Claim 1. $\mathcal{F}'$ is a baled pre-tree, i.e. a partial order such that the quotient forms a partial order with a maximal element and without the maximal element the partial order is a tree.

Proof of Claim 1. $\mathcal{F}'/\equiv$ certainly has a maximal element, namely $[z]$. Now $\mathcal{F}'/\equiv \setminus\{[z]\}$ is a partial order because both $\mathcal{F}/\equiv$ and the order on T are reflexive and transitive.

To see that $\mathcal{F}'/\equiv \setminus\{[z]\}$ forms a tree it remains to show that the predecessors of any element of $\mathcal{W}' \setminus \{[z]\}$ are linearly ordered in $\mathcal{R}'$. So let $(x,t) \in \mathcal{W}' \setminus \{[z]\}$ and $Y = \{(y,s) \mid (y,s)\mathcal{R}'(x,t)\}$ be the set of all predecessors. For any two elements $(y,s),(y',s') \in Y$ we know that $s \subseteq s'$ or $s' \subseteq s$ because T forms a tree. Say $s \subseteq s'$. Then $[y]$ is a node in the path from $[w_0]$ to $[y']$ and hence

[14] A cluster is a set C of one or more elements of $\mathcal{F}$ such that for all $a, b \in C : a \leq_{\mathcal{F}} b \leq_{\mathcal{F}} a$, i.e. all elements in C are equivalent under $\equiv$.

$y \leq y'$. So Y is linearly ordered. $\square$(Claim 1)

Since every finite baled pre-tree is a finite pre-lattice the proof is complete. $\square$

Chapter 2

Forcing

This chapter will demonstrate a rough outline of how we are going to argue when we link modal logic and forcing together. We will then prove some facts about very specific forcings and show that they behave in a suitable way. Once this is done we are ready to prove our main results in the following chapters.
Throughout this chapter it is assumed that the reader is familiar with the basic concepts of forcing as they can be found in [10].
Notations are as usual, i.e. $\mathcal{M}$ marks a countable transitive model of ZFC, $\mathbb{P}$ and $\mathbb{Q}$ are partial orders, forcing extensions of $\mathcal{M}$ are called $\mathcal{M}[G]$ where G is a generic filter for some partial order. τ and σ will be used to denote names unless otherwise defined. κ, λ are cardinals and α, β mark ordinals. Other greek letters may occure and it will follow from the context what they mean.

2.1 Motivation

Let $\Diamond\phi$ mean "there is a forcing extension where ϕ holds" and $\Box\phi$ abbreviate "ϕ holds in every forcing extension". Then the diamond can be replaced by "there is a partial order $\mathbb{P}$ and an element $p \in \mathbb{P}$ such that p forces ϕ". This can of course be expressed in the language of set theory, i.e. $\Diamond\phi$ is a sentence of the language of set theory. The same treatment goes for $\Box\phi$, so under this interpretation of the modal operators every model of set theory knows how to understand $\Diamond\phi$ and $\Box\phi$.
We will see more on this in the next chapter, and postpone a more precise treatment of this matter until then. For now we just need the forcing-interpretation of the modal operators.
We are now ready to give an important definition. This definition is the main idea in our proof that the modal logic of forcing is S4.2, as we will see in the next chapter.

Definition 2.1.1 A *button* is a statement ϕ such that $\Box\Diamond\Box\phi$ holds. A button is called *unpushed in* $\mathcal{M}$ if furthermore $\neg\phi$ holds in $\mathcal{M}$. A *switch* is a statement

ϕ such that $\Box\Diamond\phi$ as well as $\Box\Diamond\neg\phi$ hold.

The idea is that a button is a statement that can be 'turned on' via forcing and stays on in every further extension. So an unpushed button can be pushed, but cannot be unpushed again. A switch on the other hand is a statement that can be turned on and off as we please.
In the above definition we omitted the specific notion of the model where ϕ is a button or switch. As we will see later, this is not needed, as a button stays a button in every forcing extension (assuming no restrictions on the class of forcings). It will also always be clear from the context where a statement is supposed to be a button or a switch.
There are some simple facts about buttons and switches that we can prove right away just to get a better understanding of this definition. First of all there are buttons and switches in set theory. Consider for example the statement "$V \neq L$". This is forceable over every model of set theory, and once it becomes true it stays true in all forcing extensions. Hence this is a button. A switch would be the CH for example, because with forcing we can always ensure that CH or ¬CH holds in the extension. But this goes even further. Not only are there buttons and switches, but every statement is either a switch, a button or the negation of a button. To see this we have to show that any statement ϕ that is neither a button nor a switch in $\mathcal{M}$ is the negation of a button. Because ϕ is not a switch there is a partial order $\mathbb{P}$ such that in $\mathcal{M}^{\mathbb{P}}$ ϕ or $\neg\phi$ is no longer forceable. By the product lemma this means that either ϕ or $\neg\phi$ is not forceable in any further extension of $\mathcal{M}^{\mathbb{P}}$. But because ϕ is not a button there is no extension such that ϕ holds in every further extension. Hence $\mathcal{M}^{\mathbb{P}} \models \Box\neg\phi$. But this means there is a forcing extension of $\mathcal{M}$ such that in every further extension $\neg\phi$ holds, i.e. $\neg\phi$ is a button (see [1]).

As we have seen in the previous chapter, S4.2 is complete with respect to the class of finite pre-lattices. In order to show that the modal logic of forcing is S4.2 we will first of all show that all of the axioms of S4.2 hold for forcing. As a second step we will then find a link between forcing and finite pre-lattices. This is where the idea of buttons and switches comes into play. We shall show that any given finite pre-lattice can be labeled by suitable statements of the language of set theory in such a way that a statement ϕ is forceable over $\mathcal{M}^{\mathbb{P}}$ if and only if ϕ was labelled to a node in the pre-lattice that is bigger in the pre-lattice sense than the node we reached by forcing with $\mathbb{P}$. These statements are the buttons and switches such that the buttons will determine which cluster we are in and the switches will tell us our precise node in the cluster. This construction only works if we have an independence of the buttons and switches, i.e. we do not want to force some statement and thereby affect the truth value of any of the other statements. So we say a collection of buttons b_n and switches s_m is

[1] We are a bit sloppy in our argumentation. We actually only have shown that $\mathcal{M} \models \Diamond\Box\neg\phi$. But this suffices as the following argument shows: Let $\mathbb{Q}$ be an arbitrary forcing. Then $\mathcal{M}^{\mathbb{Q}\times\mathbb{P}} = \mathcal{M}^{\mathbb{P}\times\mathbb{Q}}$ again by the product lemma and is hence an extension of $\mathcal{M}^{\mathbb{P}}$, i.e. $\neg\phi$ holds in this extension. Also every further extension of $\mathcal{M}^{\mathbb{Q}\times\mathbb{P}}$ is an extension of $\mathcal{M}^{\mathbb{P}}$, i.e. $\mathcal{M}^{\mathbb{Q}\times\mathbb{P}} \models \Box\neg\phi$. Because $\mathbb{Q}$ was arbitrary $\mathcal{M} \models \Box\Diamond\Box\neg\phi$ and hence our proof is complete.

independent if all the buttons are unpushed in the ground model $\mathcal{M}$ and for any extension $\mathcal{M}^{\mathbb{P}}$ that models some pattern of the b_n, s_m any button can be turned on by forcing to some $\mathcal{M}^{\mathbb{P}\times\mathbb{Q}}$ without affecting the truth value of any of the other buttons and switches and any switch can be turned on or off by forcing to some $\mathcal{M}^{\mathbb{P}\times\mathbb{R}}$ without affecting the truth value of any of the other buttons and switches. We give a formal definition of this idea:

Definition 2.1.2 Let I, J be arbitrary subsets of ω and let $A \subseteq I$ and $B \subseteq J$. Let b_n, $n \in I$, be a collection of buttons and s_m, $m \in J$, a collection of switches. Define

$$\Theta_{A,B} = (\bigwedge_{i\in A} \Box b_i) \wedge (\bigwedge_{i\notin A} \neg\Box b_i) \wedge (\bigwedge_{j\in B} s_j) \wedge (\bigwedge_{j\notin B} \neg s_j)$$

which states that the pattern of buttons and switches that holds is specified by A and B. The family $\{b_i\}_{i\in I} \cup \{s_j\}_{j\in J}$ is called independent in $\mathcal{M}$ if

$$\mathcal{M} \models (\bigwedge_{i\in I} \neg\Box b_i) \wedge (\bigwedge_{\substack{A\subseteq I\\ B\subseteq J}} \Box(\Theta_{A,B} \rightarrow \bigwedge_{\substack{A\subseteq A'\subseteq I\\ B'\subset J}} \Diamond\Theta_{A',B'}))$$

So what we have to find is a sufficiently large collection of independent buttons and switches to label any finite pre-lattice. This means we will have to show that there is an arbitrary large family of independent buttons and switches. The rest of this chapter is devoted to prove that in fact there is such a family.

2.2 The Switches

In this section we will prove that $\phi_n \equiv$ "The GCH holds at $\aleph_{\omega+2n}$" is a switch. Furthermore we will see that this switch can be turned on and off by ω-closed forcing and is therefore a switch for ω-closed forcing. We will not use this fact until chapter 4 where we will show that the modal logic of ω-closed forcing equals S4.2.
We will see some very basic definitions and lemmata at the beginning of the next subsection. This is mainly to avoid any confusion in notations.
When we prove that the modal logic of forcing is S4.2 we will start in L. Notice that $L \models GCH$, i.e. all the switches are turned on in L. This is why we start with the discussion on how to turn them off. However, turning off these switches is possible in arbitrary models of set theory and the following takes this into account.

2.2.1 $\neg$ GCH at λ

In this subsection we will show that there is an ω-closed forcing that turns the switch off.

The following definition is very basic and it is only included to make clear that we shall distinguish between λ-closed and $< \lambda$-closed forcings as opposed to [8].

Definition 2.2.1 A partial order $\mathbb{P}$ is called λ-closed if whenever $\gamma \leq \lambda$ for every descending chain $\{p_\xi \mid \xi < \gamma\}$ of elements of $\mathbb{P}$ there is a master condition $q \in \mathbb{P}$ such that $q \leq p_\xi$ for every $\xi < \gamma$. $\mathbb{P}$ is called $< \lambda$-closed if the above holds for every γ strictly less than λ.

We shall now prove some facts about $< \lambda$-closed forcings that we will need later during our construction.

The following theorem is a tool for working with $< \lambda$-closed forcings and we will use it in many proofs in this section.

Theorem 2.2.2 *Assume* $\mathbb{P} \in \mathcal{M}$*,* $A, B \in \mathcal{M}$*,* $(\lambda$ *is a cardinal*$)^{\mathcal{M}}$*,* $(\mathbb{P}$ *is* $< \lambda$*-closed*$)^{\mathcal{M}}$ *and* $(|A| < \lambda)^{\mathcal{M}}$*. Let* G *be a* $\mathbb{P}$*-generic filter over* $\mathcal{M}$ *and let* $f \in \mathcal{M}[G]$ *with* $f : A \to B$*. Then* $f \in \mathcal{M}$*.*

Proof. Without loss of generality we can restrict ourselves to the case $A = \alpha$ for some ordinal $\alpha < \lambda$, for we can then prove the general case by letting $j \in \mathcal{M}$ be a 1-1 map from $\alpha = |A|^{\mathcal{M}} < \lambda$ onto A and apply the special case with $f \circ j : \alpha \to B$ to show that $f \circ j$, and hence f, is in $\mathcal{M}$.
Now fix $A = \alpha$ and B and let $K = ({}^\alpha B)^{\mathcal{M}} = {}^\alpha B \cap \mathcal{M}$. Let $f \in {}^\alpha B \cap \mathcal{M}[G]$. We need to show that $f \in K$.
Suppose not. We can then fix a name $\tau^G \in \mathcal{M}^{\mathbb{P}}$ for f and a $p \in G$ such that

$(*)$ $p \Vdash (\tau$ is a function from $\check{\alpha}$ to $\check{B}$ and $\tau \notin \check{K})$.

We will now construct a sequence of elements of G such that each element of the sequence will define the function τ a bit further. In the end we will be able to construct from our sequence a new function g in $\mathcal{M}$ such that g will be in K, but not its name, a contradiction.
Consider the following sequences: $\{p_\eta \mid \eta \leq \alpha\}$ in $\mathbb{P}$ and $\{z_\eta \mid \eta < \alpha\}$ in B such that

(i) $p_0 = p$

(ii) $p_\eta \leq p_\xi$ for all $\xi \leq \eta$

(iii) $p_{\eta+1} \Vdash \tau(\check{\eta}) = \check{z}_\eta$.

These sequences can be chosen in the following way:
For successor steps we are given p_η and we want to find $p_{\eta+1}$ and z_η. Since $p_\eta \leq p$ we have

$$p_\eta \Vdash \tau \text{ is a function from } \check{\alpha} \text{ to } \check{B}$$

so $p_\eta \Vdash \exists x \in \check{B}(\tau(\check{\eta}) = x)$, since a consequence of a forced statement is forced. By definition of $\Vdash$ there is an $a \in B$ such that $(\tau^G(\eta) = a)^{\mathcal{M}[G]}$. Every statement true in $\mathcal{M}[G]$ is forced by some element in G, so there is a $r \in G$ such that $r \Vdash \tau^G(\check{\eta}) = \check{a}$. Let $q \in G$ be a common extension of r and p_η. Then

$$q \Vdash \text{"}\tau \text{ is a function from } \check{\alpha} \text{ into } \check{B} \text{ and } \tau(\check{\eta}) = \check{a}\text{"}$$

Then set $p_{\eta+1} = q$ and $z_\eta = a$.
At limit steps, p_η, for η limit, may be chosen to satisfy (ii) because $\mathbb{P}$ is $< \lambda$-closed.
Now, in $\mathcal{M}$, define a function g with domain α such that $g(\eta) = z_\eta$ for each $\eta < \alpha$. Then certainly $g \in K$. Now let H be a $\mathbb{P}$-generic filter over $\mathcal{M}$ with $p_\alpha \in H$. By the closure properties of a filter, every $p_\eta \in H$, so $\tau^H(\eta) = z_\eta$ for each $\eta < \alpha$. Hence $\tau^H = g \in K$. But $p_0 = p \Vdash \tau \notin \check{K}$, so $\tau^H \notin K$. Contradiction! □

As we want to show that our switches are independent we will have to show that our forcing does not destroy the GCH at any other point. To do so we shall first show that $< \lambda$-closed forcings preseve cofinalities and cardinals $\leq \lambda$. The next lemma displays that in order to show that $\mathbb{P}$ preserves cardinals it suffices to prove that regularity is preserved in $\mathcal{M}[G]$.

Lemma 2.2.3 *Assume that $\mathbb{P} \in \mathcal{M}$, λ infinite cardinal of $\mathcal{M}$, $\mathbb{P}$ is $< \lambda$-closed and whenever κ is a regular cardinal of $\mathcal{M}$, $\kappa < \lambda$, and G is a $\mathbb{P}$ generic filter over $\mathcal{M}$, then $(\kappa$ is regular$)^{\mathcal{M}[G]}$. It then follows that $\mathbb{P}$ preserves cofinalities less than λ.*

Proof. Let γ be a limit ordinal in $\mathcal{M}$ such that, in $\mathcal{M}$, $\mathrm{cf}(\gamma) = \kappa$ for some cardinal $\kappa \leq \lambda$. Then κ is regular in $\mathcal{M}$. So by assumption κ is regular in $\mathcal{M}[G]$.
In $\mathcal{M}[G]$ let f be a strictly increasing cofinal map from κ into γ. Then $f \in \mathcal{M}$ by the above theorem and, because f is a strictly increasing cofinal map, $(\mathrm{cf}(\gamma) = \kappa)^{\mathcal{M}[G]}$. □

Lemma 2.2.4 *If $\mathbb{P}$ preserves cofinalities less than λ then $\mathbb{P}$ preserves cardinals less than λ.*

Proof. Let $\mathbb{P}$ preserve cofinalities less than λ and let α be a regular cardinal of $\mathcal{M}$ with $\omega \leq \alpha \leq \lambda$. Then by assumption $\mathrm{cf}(\alpha)^{\mathcal{M}[G]} = \mathrm{cf}(\alpha)^{\mathcal{M}} = \alpha$, i.e. α is a regular cardinal in $\mathcal{M}[G]$. For a limit cardinal $\beta \leq \lambda$ the regular cardinals $< \beta$ in $\mathcal{M}$ are unbounded in β. Since these cardinals remain regular in $\mathcal{M}[G]$, β is a limit cardinal in $\mathcal{M}[G]$ as well. Because every cardinal is either a regular or a limit cardinal (or both) the claim follows. □

Corollary 2.2.5 *Assume $\mathbb{P} \in \mathcal{M}$, $(\lambda$ is a cardinal$)^{\mathcal{M}}$ and $(\mathbb{P}$ is $< \lambda$-closed$)^{\mathcal{M}}$. Then $\mathbb{P}$ preserves cofinalities $\leq \lambda$ (and hence, by 2.2.4, cardinals $\leq \lambda$).*

Proof. Suppose not. We then find a $\kappa \leq \lambda$ such that $(\kappa$ is regular$)^{\mathcal{M}}$ but $(\kappa$ is singular$)^{\mathcal{M}[G]}$ using 2.2.3. Thus, in $\mathcal{M}[G]$, there is a cofinal function $f : \alpha \to \kappa$ for some $\alpha < \kappa \leq \lambda$. So by 2.2.2 $f \in \mathcal{M}$, which means that f is a cofinal function from α to κ in $\mathcal{M}$, a contradiction! □

We shall start with the general definition of our forcing and prove the relevant facts before stating the precise partial order that is needed to turn the switch off.

Definition 2.2.6 Let us denote the forcing of partial functions from some set I into some set J by $Fn(I, J, \lambda)$, where every function in $Fn(I, J, \lambda)$ has cardinality less than λ (see [2]). As usual, the conditions of this forcing are ordered by end-extension, so

$$Fn(I, J, \lambda) = \{p \mid p \text{ is a function } \wedge dom(p) \subseteq I \wedge ran(p) \subseteq J \wedge |p| < \lambda\}$$

and $p \leq q$ iff $p \supseteq q$.

Since we now have established an easy way to prove that cardinals $\leq \lambda$ are preserved, we just need to show that Fn is $< \lambda$-closed. This is not always, but in many cases, true, as the following lemma ensures.

Lemma 2.2.7 *If λ is regular, then $Fn(I, J, \lambda)$ is $< \lambda$-closed.*

Proof. Given a decreasing sequence $\{p_\xi \mid \xi < \gamma\}$ of elements of $Fn(I, J, \lambda)$ for some $\gamma < \lambda$, set $q = \bigcup\{p_\xi \mid \xi < \gamma\}$. Then $q \leq p_\xi \forall \xi < \gamma$, $|q| < \lambda$ since $|p_\xi| < \lambda$ and λ is regular. □

Corollary 2.2.8 *If λ is regular, then $Fn(I, J, \lambda)$ preserves cardinals $\leq (\lambda)^{\mathcal{M}}$*

Proof. Use 2.2.7 for $< \lambda$-closedness and then apply 2.2.5. □

We now want to see that Fn-forcing may destroy the GCH at regular cardinals of the form $\lambda = \aleph_{\omega+2n+1}$. Actually it destroys the GCH for arbitrary cardinals. But what we now need to verify is that Fn-forcing is precise enough to only change the cardinality of the power set at $\aleph_{\omega+2n+1}$ and does not interfere with cardinals of the form $\aleph_{\omega+2m+1}$ for $m \neq n$. We will need this precision when we want to show independence of the switches.

Theorem 2.2.9 *Fix some $n < \omega$ and let $\lambda = \aleph_{\omega+2n+1}$. Let $\mathcal{M} \models 2^\lambda = \lambda^+$ and set $\mathbb{P} = Fn(\lambda^{++} \times \lambda, 2, \lambda)$ and let G be a $\mathbb{P}$-generic filter over $\mathcal{M}$. Then $\mathcal{M}[G] \models 2^{\aleph_{\omega+2n+1}} > \aleph_{\omega+2n+2}$. Furthermore it follows that for all $m < \omega$ with $m \neq n$ if $\mathcal{M} \models 2^{\aleph_{\omega+2m+1}} = \aleph_{\omega+2m+2}$ then $\mathcal{M}[G] \models 2^{\aleph_{\omega+2m+1}} = \aleph_{\omega+2m+2}$.*

Proof. Note that λ is a successor cardinal and hence regular. Therefore Theorem 2.2.2 applies and yields $({}^\theta 2)^{\mathcal{M}} = ({}^\theta 2)^{\mathcal{M}[G]}$ for all $\theta < \lambda$ by $< \lambda$-closedness. The claim follows for all $m < n$.

In order to show $(2^\lambda \geq \lambda^{++})^{\mathcal{M}[G]}$ we use a density argument. Since G is a generic filter we know that $\bigcup G$ is a function from $\lambda^{++} \times \lambda$ into 2. So we can obtain a λ^{++}-sequence $(f_\alpha \mid \alpha < \lambda^{++})$ of functions from λ into 2 by letting

[2] This notation is taken from Kunen[8]

$$f_\alpha(\delta) = (\bigcup G)(\alpha, \delta).$$

This sequence is in $\mathcal{M}[G]$ by the preservation of cardinals below λ and the fact that $G \in \mathcal{M}[G]$. To prove that all the f_α are distinct consider

$$D_{\alpha,\beta} = \{p \in \mathbb{P} \mid \exists \theta < \lambda((\alpha,\theta) \in dom(p) \wedge (\beta,\theta) \in dom(p) \wedge p(\alpha,\theta) \neq p(\beta,\theta)\}.$$

This set is dense because for every $p \in \mathbb{P}$ that is not in $D_{\alpha,\beta}$ one can extend p to some q such that $q \in D_{\alpha,\beta}$: just let $q \restriction dom(p) = p$ and pick some θ such that (α, θ) and (β, θ) are not in $dom(p)$ and let $q(\alpha, \theta) = 0$ and $q(\beta, \theta) = 1$. Such α, β can be found because the elements of $\mathbb{P}$ have cardinality less than λ.
Using genericity of G it is now clear that $(\bigcup G)(\alpha, \delta) \neq (\bigcup G)(\beta, \delta)$ for all $\alpha, \beta < \lambda^{++}$, $\alpha \neq \beta$. So all the f_α's are distinct and hence $(2^\lambda \geq \lambda^{++})^{\mathcal{M}[G]}$.

It remains to show the claim for $m > n$. We will show the stronger statement $(2^\theta = \kappa)^{\mathcal{M}} \implies (2^\theta = \kappa)^{\mathcal{M}[G]}$ for all $\theta \geq \lambda^{++}$. This is done by a nice-name argument.

Fix some arbitrary $(\theta \geq \lambda^{++})^{\mathcal{M}}$ and let $\tau^G \subset \theta$ be a subset of θ in $\mathcal{M}[G]$. Suppose σ is a nice-name for τ^G, so

$$\sigma = \{((\alpha, i_{p_\alpha})\check{}, p_\alpha) \mid p_\alpha \in A_\alpha, i_{p_\alpha} = \begin{cases} 1 & \text{if } p_\alpha \Vdash \alpha \in \tau \\ 0 & \text{if } p_\alpha \Vdash \alpha \notin \tau \end{cases}\}$$

where the A_α's are maximal antichains in $\mathbb{P}$ and $(\alpha < \theta)^{\mathcal{M}}$. We want to see that there are many such nice-names. In order to do so we need the following claim.
Claim 1. $\mathbb{P}$ has the $(2^{<\lambda})^+$-c.c.
Proof of Claim 1.

Let $\chi = (2^{<\lambda})^+$ and suppose $\{p_\xi \mid \xi < \chi\}$ forms an antichain in $\mathbb{P}$. Because λ is regular $(2^{<\lambda})^{<\lambda} = 2^{<\lambda}$, so $\forall \alpha < \chi : (\left|\alpha^{<\lambda}\right| \leq \left|(2^{<\lambda})^{<\lambda}\right| = \left|2^{<\lambda}\right| < \chi)$.
We now apply the Δ-system-lemma to obtain a set $A \subset \chi$ with cardinality χ such that $\{dom(p_\xi) \mid \xi \in A\}$ forms a Δ-system with some root r. Because $p_\xi \in Fn(\lambda^{++} \times \lambda, 2, \lambda)$ we have $\overline{\overline{p_\xi}} < \lambda$ and hence $\overline{\overline{r}} < \lambda$. But then $2^r < \chi$ by the above argument, so there are less than χ possibilities for $p_\xi \restriction r$. So we find a $B \subset A$, $\overline{\overline{B}} = \chi$, such that $p_\xi \restriction r = p_{\xi'} \restriction r$ for every $\xi, \xi' \in B$. So all those $p_\xi, p_{\xi'}$ are compatible, thus the $\{p_\xi \mid \xi < \chi\}$ did not form an antichain. □(Claim 1)

So every antichain A_α that occurs in the nice-name has cardinality $< (2^{<\lambda})^+ \leq (2^\lambda)^+ = \lambda^{++}$ in $\mathcal{M}$ by $\mathcal{M} \models 2^\lambda = \lambda^+$.

Still working in $\mathcal{M}$ the cardinality of $\mathbb{P}$ equals $(\lambda^{++} \cdot \lambda)^{<\lambda} = (\lambda^{++})^{<\lambda} = \lambda^{++} \cdot \lambda^{<\lambda} = \lambda^{++}$ by $(2^\lambda = \lambda^+)^{\mathcal{M}}$ and Hausdorff. Therefore there are at most $(2^{\lambda^{++}})^{\lambda^{++}}$ many antichains in $\mathcal{M}$ and hence at most $((2^{\lambda^{++}})^{\lambda^{++}})^\theta = 2^{\lambda^{++} \cdot \lambda^{++} \cdot \theta} = 2^\theta$ many nice-names for θ in $\mathcal{M}$. This proves the theorem. □

2.2.2 GCH at λ

Now that we have seen that the GCH can be turned off at any given cardinal λ it remains to show that $2^\lambda = \lambda^+$ is also forceable in any model of set theory. The forcing used will again be ω-closed and will behave nicely, which will allow us to show independence of the switches in the next section. ω-closedness is not needed until chapter 4.

Definition 2.2.10 Suppose that κ, λ are cardinals. The collapse-forcing is a forcing, i.e., a partial order, of the form

$$Col(\kappa, \lambda) = \{p \mid \exists \xi \in Ord \text{ less than } \kappa \ (p : \xi \to \lambda)\} \text{ (see [3])}$$

ordered by end-extension, i.e. $p \leq q$ iff $p \upharpoonright dom(q) = q$.

Suppose that $(\lambda < \kappa)^{\mathcal{M}}$. Now the union over any $Col(\kappa, \lambda)$-generic filter G is a surjection from κ onto λ (see proof of 2.2.13). But such a surjection is already in $\mathcal{M}$ by the assumption $(\lambda < \kappa)^{\mathcal{M}}$. The extension may include some Cohen generics but we are not interessted in this matter. We shall therefore always implicitly assume that $(\kappa < \lambda)^{\mathcal{M}}$ when forcing with $Col(\kappa, \lambda)$.

Lemma 2.2.11 *Let κ be a regular cardinal and λ be a cardinal. Then $Col(\kappa, \lambda)$ has the $(\lambda^{<\kappa})^+$-c.c. and is $<\kappa$-closed.*

Proof. To prove the $(\lambda^{<\kappa})^+$-c.c. bear in mind that $\overline{\overline{\mathbb{P}}} = \lambda^{<\kappa}$. But every antichain is at most as big as the partial order. So the claim follows.
To show that $Col(\kappa, \lambda)$ is $<\kappa$-closed let $\{p_\xi \mid \xi < \delta\}$ be a descending chain in $Col(\kappa, \lambda)$ with $p_\xi : \xi \to \lambda$ where $\delta < \kappa$. Then $|p_\xi| = \overline{\overline{\xi}}$ so $\bigcup\{p_\xi \mid \xi < \delta\}$ is a union of δ-many elements, each smaller than κ. So $\overline{\overline{\bigcup\{p_\xi \mid \xi < \delta\}}} < \kappa$ because κ is regular. Hence $\bigcup\{p_\xi \mid \xi < \delta\} \in Col(\kappa, \lambda)$. Trivially $\bigcup\{p_\xi \mid \xi < \delta\} \leq_{\mathbb{P}} p_\xi$ holds for every $\xi < \delta$. □

Corollary 2.2.12 *Suppose that, in $\mathcal{M}$, κ is a regular cardinal and λ is a cardinal. Then $Col(\kappa, \lambda)$ preserves cardinals $\leq \kappa$ and $\geq ((\lambda^{<\kappa})^+)^{\mathcal{M}}$*

Proof. Preservation of cardinals $\leq \kappa$ is due to the fact that $Col(\kappa, \lambda)$ is $<\kappa$-closed (see Lemma 2.2.11) For cardinals $\geq ((\lambda^{<\kappa})^+)^{\mathcal{M}}$ use the $\lambda^{<\kappa}$-c.c. (again, Lemma 2.2.11). □

Theorem 2.2.13 *Suppose that κ and λ are infinite cardinals, $\mathcal{M} \models \kappa < \lambda$ and take $\mathbb{P}$ to be $Col(\kappa, \lambda)$. For any $\mathbb{P}$-generic filter G over $\mathcal{M}$ we then have: $\mathcal{M}[G] \models \kappa = \overline{\overline{\lambda}}$.*

[3] Note that $Col(\kappa, \lambda)$-forcing is basically $Fn(\kappa, \lambda, \kappa)$-forcing. We introduced this forcing-notion for two reasons. First of all the Col-forcing is well known and the Col-notation is usual. Secondly it makes understanding easier to insert a different name for the forcing that turns the switch back on again.

Proof. In $\mathcal{M}[G]$ we need to find a surjection from κ onto λ. As mentioned above, $\bigcup G$ is the needed function. To prove this set $F = \bigcup G$ and

$$D_\xi = \{p \mid \xi \in dom(p)\}$$

Now D_ξ is dense in $\mathbb{P}$ since if $\xi \notin dom(q)$, then p with $p \upharpoonright dom(q) = q$ and $p(\xi) = 0$ is stronger than q in $\mathbb{P}$ and $p \in D_\xi$. Hence, F is a well-defined function.
To show that F is onto, suppose $\delta < \lambda$ and let

$$D_\delta = \{p \mid \delta \in ran(p)\}$$

For any $q \notin D_\delta$ there is a $p \in D_\delta$ with $p < q$. This is because κ is an infinite cardinal, so κ is a limit ordinal. Therefore there is some $\xi < \kappa$ with $\xi \notin dom(q)$. Then simply set $p \upharpoonright dom(q) = q$ and $p(\xi) = \delta$.
We have shown that D_δ is dense for all $\delta < \lambda$. That means that for all $\delta < \lambda$ there is a $\xi < \kappa$ and a $p \in G$ such that $(p(\xi) = \delta)$, i.e., there is some $\xi < \kappa$ with $(F(\xi) = \delta)$, so F is a surjection from κ onto λ.

The proof is completed as the injection from κ into λ witnessing $\mathcal{M} \models \kappa \leq \overline{\overline{\lambda}}$ is still an element of the extension. □

2.3 The Buttons

In this section we shall see that, in L, there are ω-many buttons. In 2.4 we will prove that these buttons and the switches defined in the previous section form an independent family. As mentioned before, this is what we need in order to show that the modal logic of forcing is included within S4.2, as will be shown in chapter 3.

Hamkins and Löwe defined (see [4])

$$\phi'_n = \text{"}\omega^L_{2n+1} \text{ is not a cardinal"}.$$

They claimed that these buttons form an independent family of buttons in any model of $V = L$. They did not prove this but referred the reader to [7] 15.21, i.e. collapse forcing. But the following shows that collapse forcing is in fact not precise enough to prove independence.

We have seen in the previous section that forcing of the form

$$\mathbb{P} = Col(\overline{\overline{\omega^L_{2n}}}, \omega^L_{2n+1})$$

makes ϕ'_n true, as it forces $\overline{\overline{\omega^L_{2n}}} = \overline{\overline{\omega^L_{2n+1}}}$ in the extension. Obviously once ϕ'_n is true in one extension it stays true in every further extension, so the ϕ'_n's are buttons and Col-forcing proves it. So problems only arise when we want to show independence by using collapse forcing.

[4] To be more precise they defined $\psi'_n \equiv \text{"}\omega^L_n$ is not a cardinal". See [5] page 1768, proof of Lemma 6.1. Such a ψ'_n generates the same results as the ϕ'_n presented here, one only has to restrict the possible n. With the ϕ'_n we work around this subtlety.

Claim 1. The use of Col-forcing to push the ϕ'_n is not sufficient when proving independency.

Proof of Claim 1.

Consider the buttons ϕ'_1 and ϕ'_2. Obviously $L \not\models \neg\phi'_1 \wedge \neg\phi'_2$. Now consider the forcing $\mathbb{P} = Fn(\omega_6 \times \omega_1, 2, \omega_1)$ and let G be $\mathbb{P}$-generic over L. By the previous results we now know that $L[G] \models 2^{\omega_1} = \omega_6$. Furthermore $\mathbb{P}$ preserves cardinals (see section 2.2), hence $L[G] \not\models \neg\phi'_1 \wedge \neg\phi'_2$. If we now try to push ϕ'_1 we would ordinarily force with $\mathbb{Q} = Col(\omega_2, \omega_3)$. But because we have set the cardinality of the power set of ω_1 "high enough" we also collapse ω_6. To see this let G be the $Col(\omega_2, \omega_3)$-generic filter. Then $\bigcup G$ is a surjection from ω_2 onto ω_3. We now define

$$F : \mathscr{P}(\omega_1) \to \omega_2$$

by

$$F(x) = \text{"The smallest } \xi \text{ such that } \bigcup G \restriction [\omega_1 \cdot \xi, \omega_1 \cdot (\xi + 1)) \text{ is characteristic function of } x\text{"}$$

To see that this function is well-defined consider, for any set $A \subset \omega_1$, the set

$$D_A = \{q \in Col(\omega_2, \omega_3) \mid \exists \xi \; q \restriction [\omega_1 \cdot \xi, \omega_1 \cdot (\xi + 1)) \text{ is characteristic function of } A\}$$

Now for every $p : \alpha \to \omega_3$ in $Col(\omega_2, \omega_3)$ we may assume $\alpha = \omega_1 \cdot \xi$ for some ξ as one could enlarge the domain if needed. We then define a p' as follows:

$$p' : \alpha + \omega_1 \to \omega_3$$

with

$$p'(\beta) = p(\beta) \text{ if } \beta < \alpha \text{ and } p'(\alpha + \beta) = \chi_A(\beta)$$

Certainly $p' \restriction dom(p) = p$, hence $p' \leq p$. But also $p' \in D_A$, which implies that D_A is dense. Therefore F is well-defined.

F is also injective because if $F(x) = F(y) = \xi$ then $\bigcup G \restriction [\omega_1 \cdot \xi, \omega_1 \cdot (\xi + 1))$ is characteristic function of x and of y which implies $x = y$.

Therefore ω_6 got collapsed and hence especially ω_5 got collapsed. But that means that ϕ'_2 was pushed when we only wanted to push ϕ'_1.

□(Claim 1)

Note that we needed the exact forcing to show that F is well-defined. We hence have not shown that ω_5 gets collapsed in every forcing extension. So it might be that there is some forcing that only pushes ϕ'_1 and not ϕ'_2. We have therefore only shown that the argument presented in [5] is incorrect but not that the family of the ϕ'_n is not-independent.

Instead of looking for a proof that only ϕ'_1 can be pushed we define new buttons which are easier to be seen independent. They will also allow us to generalise the proof to ω-closed forcings, and, for some κ, $< \kappa$-closed forcings.

The main idea is the following:

We still want to use the idea that "ω_m^L is not a cardinal" is a button. But we modify it to ensure that the GCH holds at ω_m^L (note that the lack of the GCH was what caused the problems with the buttons as defined by Hamkins and Löwe in [5]). Also it turns out that the independence proof becomes much easier once we ensure that the next two successor cardinals of ω_m^L are still cardinals, which is why we define

$$\phi_n \equiv \text{"}\omega_{3n}^L \notin Card \vee \omega_{3n+1}^L \notin Card \vee \omega_{3n+2}^L \notin Card \vee \left|\mathscr{P}(\omega_{3n}^L)\right| > \left|\omega_{3n+1}^L\right|\text{"}$$

To see that ϕ_n is a button note that if a model $\mathcal{M}$ models $\neg\phi_n$ then $Col(\omega_{3n}^L, \omega_{3n+2}^L)$ computed in $\mathcal{M}$ ensures $\Box\phi_n$ in every further extension of $\mathcal{M}$, since a collapsed cardinal can not be "uncollapsed" by means of forcing.

2.4 Independence

This section is divided into three parts: the first is devoted to the proof that the switches we defined form an independent family in L.

We first consider the case of a forcing extension $\mathcal{M}$ of L where the GCH fails at some cardinals, and we want to recover it at one explictit cardinal while not affecting the cardinality of the powerset of any other cardinal. There are some subtleties to be taken care of and a technical treatment seemed unavoidable. The idea is: because we need the GCH at some points for nice chain-condition-properties we start by forcing it wherever we need it. Then we destroy it again at all the points we did not want to obtain it. Once all this is done we code the forcings used into one forcing using the concept of iterated forcing. This will have made sure that the GCH was only destroyed at one point of interest.

To force ¬GCH at a certain cardinal does not involve any difficulties.

The second part shows that the buttons are independent. We will also see a brief note on what happens when we do not incorporate "ω_{3n+2}^L is not a cardinal" into the buttons.

The last part then shows that the buttons and switches are independent from one another. This is merely uses closedness and chain condition properties.

Note that we only show that an arbitrary large but finite family of these buttons and switches is independent, while our definition of independence did not request the family to be finite. However, finite families suffice by the finite frame property of S4.2, as will be seen in the next chapter.

2.4.1 Independence of the switches

Let

$$\psi_n \equiv 2^{\kappa_n} = \kappa_n^+$$

where

$$\kappa_n = \aleph_{\omega+2n+1}.$$

We have already seen that the ψ_n are switches and now we want to show that they can be turned on and off individually.

Let $A \subset \omega$ be finite. The goal of this subsection is to show that $\{\psi_n \mid n \in A\}$ is an independent family.

Let $\mathcal{M} = L[H]$ be a forcing extension of L and suppose

$$\mathcal{M} \models \bigwedge_{n\in A'} \psi_n \wedge \bigwedge_{n\notin A'} \neg\psi_n \text{ for some } A' \subset A \text{ (see }^5)$$

We need to see that there is a partial order $\mathbb{P}$ and a generic filter G such that

$$(\bigstar)\ \mathcal{M}[G] \models \bigwedge_{n\in A''} \psi_n \wedge \bigwedge_{n\notin A''} \neg\psi_n \text{ for any } A'' \subset A.$$

Since A is finite it suffices to show that

(i) there is a partial order $\mathbb{P}$ and a generic filter G with

$$\mathcal{M}[G] \models \bigwedge_{n\in A'\cup\{m\}} \psi_n \wedge \bigwedge_{n\notin A'\cup\{m\}} \neg\psi_n \text{ for all } m \in A$$

and

(ii) there is a partial order $\overline{\mathbb{P}}$ and a generic filter $\overline{G}$ with

$$\mathcal{M}[\overline{G}] \models \bigwedge_{n\in A'\setminus\{m\}} \psi_n \wedge \bigwedge_{n\notin A'\setminus\{m\}} \neg\psi_n \text{ for all } m \in A.$$

If (i) and (ii) hold, then $(\bigstar)$ can be achieved by iterated forcing.

We will first show (i).

Note that in $\mathcal{M}$ ψ_m can be forced by

$$\mathbb{P}^* = Col(\kappa_m^+, 2^{\kappa_m}).$$

Because $\mathbb{P}^*$ is $< \kappa_m^+$-closed we have

$$({}^A B)^{\mathcal{M}} = ({}^A B)^{\mathcal{M}[G^*]}$$

if G^* is $\mathbb{P}^*$-generic over $\mathcal{M}$ and $|A| < \kappa_m^+$. Hence $\mathcal{M} \models \psi_n$ iff $\mathcal{M}[G^*] \models \psi_n$ for all $n < m$.

But $\mathbb{P}^*$ does not suffice as the following shows.

$\mathbb{P}^*$ has the $((2^{\kappa_m})^{<\kappa_m^+})^+ = (2^{\kappa_m})^+$-c.c. Suppose there is a l such that $m < l$, $l \in A$, $\mathcal{M} \models \psi_l$ and $\mathcal{M} \models \kappa_l < 2^{\kappa_m}$. Such l may exist, because $\mathcal{M}$ was an arbitrary forcing extension of L. Then there are at most $(|\mathbb{P}^*|^{(2^{\kappa_m})})^{\kappa_l} = ((2^{\kappa_m})^{(2^{\kappa_m})})^{\kappa_l} = 2^{(2^{\kappa_m})}$- many nice-names for subsets of κ_l. But we want independence, i.e. we need to show $\mathcal{M}[G^*] \models \psi_l$. By $((2^{\kappa_m}) > \kappa_l)^{\mathcal{M}}$ we see that the upper bound for subsets of the κ_l is too big, i.e. κ_l got collapsed and we therefore affected the GCH at κ_l.

We work around this problem by using iterated forcing. The idea is to ensure the GCH at all important stages, i.e. at every κ_l with $l \in A$, starting with the

[5] $n \notin A'$ is taken to mean $n \in A \setminus A'$.

GCH at κ_m. Once this is established we destroy the GCH again at every κ_l where $\mathcal{M} \models \neg\psi_l$.

To show that this works let

- $n_0 = m$
- $n_{i+1} = min\{n \in A \setminus \{n_0, ..., n_i\} \mid m \leq n\}$

End the sequence of the n_i if there is no $n' \in \{n \in A \setminus \{n_o, ..., n_i\}\}$ with $m \leq n'$.

Now let $\mathbb{P}_{n_i} = Col(\kappa_{n_i}^+, 2^{\kappa_{n_i}})$. Note that we do not want to calculate these forcings in $\mathcal{M}$ but rather want to use a short notation of the forcings which we are going to use when we calculate these forcings in the extensions.

Let G_{n_0} be $\mathbb{P}_{n_0}$-generic over $\mathcal{M}$. Then

$$\mathcal{M} \models \psi_n \text{ iff } \mathcal{M}[G_{n_0}] \models \psi_n \text{ for all } n < n_0,\, n \in A$$

by $< \kappa_{n_0}^+$-closedness (see above). Furthermore $\mathcal{M}[G_{n_0}] \models \psi_{n_0}$. Now construct $\mathbb{P}_{n_1}$ in $\mathcal{M}[G_{n_0}]$ to arrive at a forcing extension $\mathcal{M}[G_{n_0}][G_{n_1}]$ where

Claim 2. $\mathcal{M}[G_{n_0}][G_{n_1}] \models \psi_{n_0} \wedge \psi_{n_1}$

Proof of Claim 2. $\mathcal{M}[G_{n_0}][G_{n_1}] \models \psi_{n_1}$ is clear by definition of the partial order.

Because $\mathbb{P}_{n_1}$ is $< \kappa_{n_1}^+$-closed we know that the GCH below $\kappa_{n_1}^+$ did not get changed. Especially $2^{\kappa_{n_0}} = \kappa_{n_0}^+$ holds in the extension, i.e. $(\psi_{n_0})^{\mathcal{M}[G_{n_0}][G_{n_1}]}$.

□(Claim 2)

We now repeat this operation finitely many times until we arrive at some partial order $\mathbb{P}_{n_k}$ and a generic G_{n_k} such that $\mathcal{M}' := \mathcal{M}[G_{n_0}], ..., [G_{n_k}] \models \psi_n$ for all $n \geq m$, $n \in A$. We have hence ensured that the GCH holds at all κ_n with $n \geq m$, $n \in A$. We now have to destroy it again at all κ_n where $\mathcal{M} \models \neg\psi_n$ and $n > m$. Note that, for obvious reasons, we leave the GCH at κ_m untouched.

Define

- $\overline{n}_0 = min\{n \in A \mid m < n \wedge \mathcal{M} \models \neg\psi_n\}$
- $\overline{n}_{i+1} = min\{n \in A \setminus \{\overline{n}_0, ..., \overline{n}_i\} \mid m < n \wedge \mathcal{M} \models \neg\psi_n\}$

and let the sequence end if there is no $n \in A \setminus \{\overline{n}_0, ..., \overline{n}_i\}$ with $m < n \wedge \mathcal{M} \models \neg\psi_n$.

Write $\mathbb{Q}_{\overline{n}_i}$ for $Fn(\kappa_{\overline{n}_i}^{++} \times \kappa_{\overline{n}_i}, 2, \kappa_{\overline{n}_i})$. Again the construction of the partial orders will take place in various models.

Let $G_{\overline{n}_0}$ be a $\mathbb{Q}_{\overline{n}_0}$-generic filter over $\mathcal{M}'$ and let $\overline{n}_0 < l$. By the $(2^{<\kappa_{\overline{n}_0}})^+ \leq (2^{\kappa_{\overline{n}_0}})^+ = \kappa_{\overline{n}_0}^{++}$-c.c. there are at most $(|\mathbb{Q}_{\overline{n}_i}|^{\kappa_{\overline{n}_0}^+})^{\kappa_l} = (((\kappa_{\overline{n}_0})^{<\kappa_{\overline{n}_0}})^{\kappa_{\overline{n}_o}^+})^{\kappa_l} \leq (((\kappa_l)^{\kappa_l})^{\kappa_l})^{\kappa_l} = 2^{\kappa_l} = \kappa_l^+$ many nice-names for subsets of κ_l in $\mathcal{M}'$, i.e.

$$\mathcal{M}'[G_{\overline{n}_0}] \models \psi_l \text{ for all } l > \overline{n}_o,\, l \in A.$$

In $\mathcal{M}'[G_{\overline{n}_0}]$ we then construct $\mathbb{Q}_{\overline{n}_1}$ and arrive at an extension $\mathcal{M}'[G_{\overline{n}_0}][G_{\overline{n}_1}]$ where $\neg\psi_{\overline{n}_0} \wedge \neg\psi_{\overline{n}_1}$ holds. Repeat this operation until $\neg\psi_{\overline{n}_i}$ holds for all $\overline{n}_i$. Say we need $\overline{k}$-many steps. Call the obtained model $\mathcal{M}''$.

Because $\mathbb{Q}_{\overline{n}_i}$ is $< \kappa_{\overline{n}_i}$-closed we do not affect any ψ_n with $n < \overline{n}_i$. We have hence obtained a model $\mathcal{M}''$ with

$$\mathcal{M} \models \psi_n \text{ iff } \mathcal{M}'' \models \psi_n \text{ for all } n \in A \setminus \{m\} \text{ and } \mathcal{M}'' \models \psi_m.$$

Therefore $\mathcal{M}''$ is a model as we wanted it to be, if it is a forcing extension of $\mathcal{M}$. But this holds by iterated forcing, i.e. $\mathcal{M}'' = \mathcal{M}[G]$ for a $\mathbb{P}$-generic filter G where

$$\mathbb{P} = \mathbb{P}_{n_0} * \dots * \mathbb{P}_{n_k} * \mathbb{Q}_{\overline{n}_0} * \dots * \mathbb{Q}_{\overline{n}_{\overline{k}}}$$

Remark 2.4.1 Note that every $\mathbb{P}_{n_i}$, $\mathbb{Q}_{\overline{n}_i}$ was ω-closed. Hence $\mathbb{P}$ is also ω-closed. We will need this fact in chapter 4.

It is now easy to see that (ii) can also be achieved. We are in a situation where ψ_m holds in $\mathcal{M}$ (i.e. (GCH at $\kappa_m)^{\mathcal{M}}$) and we need to see that we can force $\neg\psi_m$. As seen in 2.2 this may be done by forcing of the form $\mathbb{Q} = Fn(\kappa_m^{++} \times \kappa_m, 2, \kappa_m)$. Because this forcing is $< \kappa_m$-closed no new subsets below κ_m are added. Hence for $n < m$

- $\mathcal{M} \models \psi_n$ iff $\mathcal{M}^{\mathbb{Q}} \models \psi_n$

Now let $m < l$. Because $\mathbb{Q}$ has the κ_m^{++}-c.c. and $(\kappa_m^{++} \leq \kappa_l)^{\mathcal{M}}$ there are at most $(|\mathbb{Q}|^{\kappa_m^{++}})^{\kappa_l} = |\mathbb{Q}|^{(\kappa_m^{++} \cdot \kappa_l)} = |\mathbb{Q}|^{\kappa_l} = |\mathbb{Q}| \cdot 2^{\kappa_l} = 2^{\kappa_l}$ many nice-names for subsets of κ_l in $\mathcal{M}$ by Hausdorff and the fact that the forcing is small relative to 2^{κ_l}. This has shown

- $\mathcal{M} \models \psi_l$ iff $\mathcal{M}^{\mathbb{Q}} \models \psi_l$

So the proof for (ii) is complete.

2.4.2 Independence of the buttons

Recall that

$$\phi_n \equiv \text{"}\omega_{3n}^L \notin Card \vee \omega_{3n+1}^L \notin Card \vee \omega_{3n+2}^L \notin Card \vee \left|\mathscr{P}(\omega_{3n}^L)\right| > \left|\omega_{3n+1}^L\right|\text{"}$$

Because we only need to show that the ϕ_n form a *finite* independent family of buttons in L we argue as we did for the independence of the switches to see that the following suffices to show:

- Let $m < n < k < \omega$ and $\mathcal{M}$ be a forcing extension of L such that $\mathcal{M} \models \neg\phi_m \wedge \neg\phi_n \wedge \neg\phi_k$. Then there is a forcing $\mathbb{P}$ such that $\mathcal{M}^{\mathbb{P}} \models \neg\phi_m \wedge \Box\phi_n \wedge \neg\phi_k$.
- $L \models \neg\phi_m \wedge \neg\phi_n \wedge \neg\phi_k$.

The second point is trivial by the GCH in L. So only the first point remains to be shown. But we can weaken it to:

- $\exists \mathbb{P} : \mathcal{M}^{\mathbb{P}} \models \neg\phi_m \wedge \phi_n \wedge \neg\phi_k$.

This is because once ϕ_n becomes true, it stays true in all further extensions. A proof of this obviously only needs to consider an extension $\mathcal{N}$ of L with $\mathcal{N} \models \omega_{3n}^L, \omega_{3n+1}^L, \omega_{3n+2}^L \in Card \wedge \left|\mathscr{P}(\omega_{3n}^L)\right| > \omega_{3n+1}^L$. Then of course $\mathcal{N} \models \phi_n$. If we were to ensure the GCH at ω_{3n}^L again we would have to collapse $\mathscr{P}(\omega_{3n}^L)$. By $\left|\mathscr{P}(\omega_{3n}^L)\right| > \omega_{3n+1}^L$ this implies that ω_{3n+2}^L would get collapsed as well and hence would not be a cardinal in the extension, i.e. ϕ_n would still be pushed in the extension.

Suppose we had not included "$\omega_{3n+2}^L \notin Card$" into the ϕ_n. Then a model $\mathcal{N}$ may exist with $\mathcal{N} \models \phi_n \wedge \neg\Box\phi_n$ because the GCH at ω_{3n}^L may be recovered without affecting ω_{3n}^L and ω_{3n+1}^L. To see that this is a problem recall the definition of independence. We needed the buttons to be pushed, i.e. $\Box\phi_m$ to hold. Therefore in our above example even though ϕ_n holds in $\mathcal{N}$ we would still have to ensure that $\Box\phi_n$ holds. We could do this by collapsing ω_{3n+1}^L onto ω_{3n}^L using Col-forcing. But suppose that $n = 1$ and $(\left|\mathscr{P}(\omega_1^L)\right| = \left|\mathscr{P}(\omega_2^L)\right| = \omega_5^L)^{\mathcal{N}}$. Then the Col-forcing is big (has cardinality ω_5^L) and, using the proof presented at the beginning of section 2.3, collapses ω_4^L. But that means that ϕ_2 would have been pushed (ω_4^L is no longer a cardinal) when we only wanted to push ϕ_1.

If we were to try and recover the GCH first and then use Col-forcing we would collapse ω_4^L while recovering the GCH.

This is the reason why we included "$\omega_{3n+2}^L \notin Card$" into the ϕ_n.

So let $\mathcal{M}$ be a forcing extension of L and $\mathcal{M} \models \neg\phi_m \wedge \neg\phi_n \wedge \neg\phi_k$ where $m < n < k < \omega$. Let $\mathbb{P} = Col(\omega_{3n}^L, \omega_{3n+1}^L)$ and G be a $\mathbb{P}$-generic filter over $\mathcal{M}$. We claim:

$$\mathcal{M}[G] \models \neg\phi_m \wedge \phi_n \wedge \neg\phi_k$$

Of course $\mathcal{M}[G] \models \phi_n$ because ω_{3n+1}^L got collapsed by $\mathbb{P}$ and is hence no longer a cardinal.

Notice that $\mathcal{M} \models \neg\phi_i$ means $(\omega_{3i}^L, \omega_{3i+1}^L, \omega_{3i+2}^L \in Card \wedge \left|\mathscr{P}(\omega_{3i}^L)\right| = \omega_{3i+1}^L)^{\mathcal{M}}$. We shall use this extensively.

By lemma 2.2.11 $\mathbb{P}$ is $< \omega_{3n}^L$-closed which implies $(\omega_{3m}^L, \omega_{3m+1}^L, \omega_{3m+2}^L \in Card)^{\mathcal{M}[G]}$ by corollary 2.2.5. Theorem 2.2.2 implies that no new subsets below ω_{3n}^L are added by $\mathbb{P}$ and hence $(\left|\mathscr{P}(\omega_{3m}^L)\right| = \omega_{3m+1}^L)^{\mathcal{M}[G]}$. This proves $\mathcal{M}[G] \models \neg\phi_m$.

We already know that $\mathbb{P}$ has the $((\omega_{3n+1}^L)^{<\omega_{3n}^L})^+$-c.c. Using Hausdorff and the GCH at ω_{3n}^L in $\mathcal{M}$ we obatain $((\omega_{3n+1}^L)^{<\omega_{3n}^L} = \omega_{3n+1}^L \cdot 2^{<\omega_{3n}^L} = \omega_{3n+1}^L)^{\mathcal{M}}$, i.e. $\mathbb{P}$ has the $(\omega_{3n+1}^L)^+$-c.c. Because ω_{3n+1}^L and ω_{3n+2}^L are both cardinals in $\mathcal{M}$ $\mathbb{P}$ has the ω_{3n+2}^L-c.c. This implies $(\omega_{3k}^L, \omega_{3k+1}^L, \omega_{3k+2}^L \in Card)^{\mathcal{M}[G]}$ by corollary 2.2.12.

To see that the GCH at ω_{3k}^L is preserved we use a nice-name argument once more:

$|\mathbb{P}| = (\omega_{3n+1}^L)^{<\omega_{3n}^L} = \omega_{3n+1}^L$ so by the $(\omega_{3n+1}^L)^+$-c.c. there are at most $(\omega_{3n+1}^L)^{\omega_{3n+1}^L}$ many antichains in $\mathbb{P}$. So there are at most $((\omega_{3n+1}^L)^{\omega_{3n+1}^L})^{\omega_{3k}^L} = (\omega_{3n+1}^L)^{\omega_{3n+1}^L \cdot \omega_{3k}^L} = (\omega_{3n+1}^L)^{\omega_{3k}^L} = 2^{\omega_{3k}^L} \cdot \omega_{3n+1}^L = 2^{\omega_{3k}^L} = \omega_{3k+1}^L$ many nice-names for substes of ω_{3k}^L in $\mathcal{M}$ and hence $\mathcal{M}[G] \models 2^{\omega_{3k}^L} = \omega_{3k+1}$, i.e. $\mathcal{M}[G] \models \neg\phi_k$.

This completes the proof of independence for the buttons.

2.4.3 Independence of the buttons and switches

The point is that we have arranged the switches "high enough" so that they do not interfere with our buttons. More precisely note that every forcing that was used to turn a switch on or off was especially $< \aleph_\omega$-closed, i.e. no new subsets below $\aleph_\omega$ were added. Hence no new models that would push the buttons were generated and therefore the buttons were not affected.

On the other hand the forcing used to push ϕ_m has the $(\overline{\overline{\omega^L_{2m+1}}})^+$-c.c., which implies that "only few" subsets above $(\overline{\overline{\omega^L_{2m+1}}})^+$ are added. Especially the GCH above $\aleph_\omega$ did not get affected.

Hence the buttons and switches are independent from one another.

Chapter 3

The Forcing Interpretation Of Modal Logic

In this chapter we will interpret the modal operator $\Diamond$ as "there is a forcing extension such that..." and the $\Box$ as "in every forcing extension...". We will then establish that, if ZFC is consistent, exactly every assertion of the modal theory S4.2 holds for the forcing interpretation of modal logic.

In his paper "A simple maximality principle" [4] Hamkins introduced the following forcing interpretation of modal logic: the statement $\Diamond\phi$ is true in some model $\mathcal{M}$ of ZFC if there is a forcing extension $\mathcal{M}^{\mathbb{P}}$ of $\mathcal{M}$ such that $\mathcal{M}^{\mathbb{P}} \models \phi$. In this case we will also call ϕ *possible* or *forceable*. The statement $\Box\psi$ then holds in $\mathcal{M}$ if for every extension $\mathcal{M}^{\mathbb{P}}$ of $\mathcal{M}$: $\mathcal{M}^{\mathbb{P}} \models \psi$. In which case ψ is called *necessary*.
The modal operators are eliminable in the language of set theory. $\Diamond\phi$ can be expressed as the statement "$\exists\mathbb{P}\ \exists p \in \mathbb{P}$ ($\mathbb{P}$ is partial order $\wedge\ p \Vdash_{\mathbb{P}} \phi$)". The statement $\Box\phi$ can be re-stated as "$\forall\mathbb{P}\ \forall p \in \mathbb{P}$ ($\mathbb{P}$ is partial order $\wedge\ p \Vdash_{\mathbb{P}} \phi$)". We will therefore use a mixed language of set theory and the modal operators, understood as abbreviations for their forcing interpretations.
Now that we have an interpretation of the modal operators, it is a natural question to ask what principles are valid for this construction. Let us write $\phi(p_1, ..., p_n)$ for a modal assertion ϕ with $p_1, ..., p_n$ occurring. Recall that a formula $\phi(p_1, ..., p_n)$ (of the modal language) is called valid on a frame if it holds under every valuation, i.e. if $\phi(p_1, ..., p_n)$ is true in every model based on that frame. Now consider a frame of models of set theory with the accessibility relation given by the forcing interpretation of the modal operators. A valuation on this frame then decides for every modal assertion $\phi(p_1, ..., p_n)$ and every node w if $\phi(p_1, ..., p_n)$ is true in w. Now w is a model of set theory and therefore for every $\psi \in \mathcal{L}_\in$ the information if ψ holds in w is already coded in w. We have seen that the modal operators are eliminable, so the question if $\phi(p_1, ..., p_n)$ holds at w depends on how the $p_1, ..., p_n$ are interpreted in the language of set

theory. We shall connect the p_i to sentences of $\mathcal{L}_\in$. Formally this is a function $\chi : \Psi \to \mathcal{L}_\in$ where Ψ is the set of propositional variables of the modal language. So $\phi(p_1, ..., p_n)$ is true in w if $\phi(\chi(p_1), ..., \chi(p_n))$ holds in w. Hence $\phi(p_1, ..., p_n)$ is valid if for every function χ as above and every node w: $w \models \phi(\chi(p_1), ..., \chi(p_n))$. We therefore define

Definition 3.0.2 A modal assertion $\phi(p_1, ..., p_n)$ is a *ZFC-provable principle of forcing* (see [1]) if for all sentences ψ_i, $0 < i \leq n$, of the language of set theory, ZFC proves every substitution instance $\phi(\psi_1, ..., \psi_n)$.

Our main goal is to show that the valid principles of forcing are excatly those of the modal theory S4.2. To do so we will prove a lemma that connects the modal language to set theory. Let us look at this lemma now:

Lemma 3.0.3 *Given some finite pre-lattice $\mathcal{F}$. Let $\mathcal{M}$ be a model based on $\mathcal{F}$ an let w_0 be a world in $\mathcal{M}$. Let $\mathcal{W}$ be a model of set theory with a sufficiently large independent family of buttons and switches. Then there is an assignment $\chi : \Psi \to \mathcal{L}_\in$ such that for every modal assertion ϕ we have*

$$(\mathcal{M}, w_0) \models \phi(q_0, ..., q_n) \text{ iff } \mathcal{W} \models \phi(\chi(q_0), ..., \chi(q_n))$$

The above lemma will be the heart of our argument for it sets an upper bound to the ZFC-provable principles of forcing. A lower bound can be easily found and it happens to equal the upper bound. In fact, the lower bound is found by just verifying that every assertion in the modal theory S4.2 holds under the forcing interpretation. We shall do so right away.

Lemma 3.0.4 *S4.2 is included within the ZFC-provable principles of forcing.*

Proof. (K)$\equiv \Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$
Suppose some model $\mathcal{M}$ where $\Box(\phi \to \psi)$ holds. If then $\Box\phi$ holds it means that in every extension ϕ holds. But by assumption $\mathcal{M} \models \Box(\phi \to \psi)$, so ψ holds in the extension by modus ponens.
(Dual) $\equiv \neg\Diamond\phi \leftrightarrow \Box\neg\phi$
From left to right note that if no extension believes ϕ, then every extension believes $\neg\phi$. On the other hand if every extension believes $\neg\phi$ then there can not be an extension believing ϕ.
(S) $\equiv \Box\phi \to \phi$
Every model is a (trivial) forcing extension of itself.
(4) $\equiv \Box\phi \to \Box\Box\phi$
A forcing extension of a forcing extension is a forcing extension of the ground

[1] We will refer to the ZFC-provable principles of forcing as the valid principles of forcing. This is somewhat informal because usually the validity depends on the valuation functions (which are functions from Φ to an element of the powerset of the worlds). But in the forcing interpretation validity depends on the functions χ from Φ to $\mathcal{L}_\in$.

model, i.e. if ϕ holds in all extensions of $\mathcal{M}$, then so does $\Box\phi$.
(.2) $\equiv \Diamond\Box\phi \to \Box\Diamond\phi$
Let ϕ be necessary in $V^{\mathbb{P}}$ and $V^{\mathbb{Q}}$ be an arbitrary extension, then ϕ holds in $V^{\mathbb{P}\times\mathbb{Q}}$ as this extends $V^{\mathbb{P}}$. So ϕ is forceable over any such extension $V^{\mathbb{Q}}$ by the Product Lemma.
Closure under modus ponens and substitution is trivial. For closure under necessitation note that if ϕ holds in all models of set theory, then ϕ holds in all extensions of any given model, i.e. $\Box\phi$ holds in all models of set theory.

□

To see explicit failing instances for any theory beyond S4.2 the reader is invited to have a look at observation 4 in [5].

Denote the ZFC-provable principles of forcing by $Force$. We have just proven that S4.2 $\subseteq Force$. To see the inverse we will use the above mentioned lemma and shall now commence the proof of it by showing some lemmata.

Lemma 3.0.5 *If $V = L$ then there are arbitrarily large (but finite) families of buttons b_i and switches s_j such that the family is independent.*

Proof. This was shown in section 2.4 □

We have seen in chapter 1 that every modal assertion ϕ not in S4.2 fails in some Kripke Model whose frame is a finite pre-lattice. We want to make use of this fact by finding a labeling of any given finite pre-lattice with statements in $\mathcal{L}_\in$ such that they imitate the modal behaviour of the structure in the forcing interpretation. We will use the concept of buttons and switches for that, so we want to find a labeling such that if a set of statements A was labeled to some world w and wRv in the lattice-order, then v was labeled by some statements B such that every statement in B is forceable over any model of set theory satisfying every statement in A. We shall use buttons to determine the cluster of the lattice and switches to determine the node within the cluster. The following lemmata show that this can be done, starting with the easier case of a finite lattice, where we only need buttons.

Lemma 3.0.6 *Let F be a finite lattice with a minimal node w_0 and $\mathcal{M}$ a model of set theory with a sufficiently large independent family of buttons b_i. Then there is an assignment that connects each node $w \in F$ to an extended Boolean combination (see [2]) p_w of the buttons such that:*

(i) In any forcing extension exactly one of the p_w is true and $\mathcal{M} \models p_{w_0}$

(ii) If $\mathcal{M}[G] \models p_w$ then p_v is forceable over $\mathcal{M}[G]$ iff $w \leq_F v$

Proof. Associate every $w \in F$ with a button b_w. For any $A \subseteq F$ let $b_A = (\bigwedge_{u\in A} \Box b_u) \wedge (\bigwedge_{u\notin A} \neg\Box b_u)$ be the sentence asserting that exactly the buttons

[2] ψ is an *extended Boolean combination* of some formulas ϕ_i if ψ is a Boolean combination of the ϕ_i or, if ψ' is an extended Boolean combination of the ϕ_i, $\psi = \Diamond\psi'$.

b_u have been pushed for $u \in A$.
(i) Set $p_w = \bigvee\{b_A \mid \bigvee(A) = w\}$ (see [3]). . Note that every subset of a finite lattice has a unique upper bound by definition. So p_w assets that some pattern A of buttons has been pushed with $\bigvee A = w$.
In any given forcing extension (of $\mathcal{M}$) exactly one pattern of buttons has been pushed. Fix some extension and call the pattern that happens to hold in this extension A. So b_A holds in this extension and therefore p_w holds for $w = \bigvee A$. Now assume that p_w holds in the extension. It then follows, because the pattern A was pushed, that $\bigvee A = w$. So p_w is true in the extension if and only if $\bigvee A = w$. Thus only one p_w can hold in any given extension.
In $\mathcal{M}$ no button has been pushed (recall that we required this for independent buttons), so $A = \emptyset$. By definition $\bigvee \emptyset =$ minimal node in F, which is w_0, so $\mathcal{M} \models p_{w_0}$.
(ii) Assume $\mathcal{M}[G] \models p_w$ and call the pattern of buttens pushed in $\mathcal{M}[G]$ A. So by (i) we know $\bigvee A = w$. Let $v \in F$ such that $w \leq_F v$. By independency of the buttons we can push the button b_v in $\mathcal{M}[G]$ without affecting the truth value of the other buttons, arriving at an extension $\mathcal{M}[G][H]$ where the pattern of buttons pushed is $A \cup \{v\}$. By $w \leq_F v$ we know $\bigvee(A \cup \{v\}) = v$, so $\mathcal{M}[G][H] \models p_v$ by (i). Therefore p_v is forceable in $\mathcal{M}[G]$ as desired.
Conversely suppose there is an extension $\mathcal{M}[G][H]$ satisfying p_v. Call the pattern of buttons pushed in this extension B. Then $\bigvee B = v$ by (i). Because $\mathcal{M}[G] \subseteq \mathcal{M}[G][H]$ and buttons can not be unpushed by forcing we know that $A \subseteq B$. Thus $\bigvee A \leq_F \bigvee B$, i.e. $w \leq_F v$. □

We shall now implement the switches to obtain the above result for (finite) pre-lattices.

Lemma 3.0.7 *Let $\mathcal{F}$ be a finite pre-lattice such that the largest cluster contains k nodes for some $k < \omega$ and let there be m clusters in $\mathcal{F}$. Let n be such that $k \leq 2^n$. Then let and $\{b_i, s_j\}_{i \leq m, j \leq n}$ be a independent family of buttons and switches in a model of set theory $\mathcal{M}$. Suppose w_0 is any node in the minimal cluster of F.*
There is an assignment from the elements of $\mathcal{F}$ to an extended Boolean combination of the buttons and switches such that if $w \mapsto p_w$

(i) In any forcing extension of $\mathcal{M}$ exactly one p_w holds and $\mathcal{M} \models p_{w_0}$.

(ii) If $\mathcal{M}[G] \models p_w$ for any G then $\mathcal{M}[G] \models \Diamond p_v$ iff $w \leq_{\mathcal{F}} v$

Proof. As said before the buttons will determine the cluster (using the above lemma) while the switches determine which node in the cluster is intended.
Denote the equivalence class of u in the quotient lattice $\mathcal{F}/\equiv$ by $[u]$. Note that $\mathcal{F}/\equiv$ is a lattice and let $p_{[u]}$ be as in the lemma above.
For a set $A \subseteq n$ let $s_A = (\bigwedge_{i \in A} s_i) \wedge (\bigwedge_{i \notin A} \neg s_i)$ be the assertion that exactly the pattern A of switches holds. Because any pattern of switches is possible by

[3] It is common practice to write $\bigvee A$ instead of $sup(A)$ when dealing with lattices.

forcing over $\mathcal{M}$ s_A is necessarily possible. Note that in any forcing extension exactly one s_A holds.
For any cluster $[u]$ let $< A_w \mid w \in [u] >$ be a partition of $\mathscr{P}(n)$, i.e. $\mathscr{P}(n) = \dot{\bigcup}_{w\in[u]} A_w$. Recall that the largest cluster contained k elements and $k \leq 2^n$, so such a partition can be found. If the pattern of switches B holds in $\mathcal{M}$, then set $A_{w_0} = \{B\}$.
Now set $s_w = \bigvee_{A\in A_w} s_A$ which asserts that the switches occur in a pattern appearing in A_w. Finally define $p_w = p_{[w]} \wedge s_w$.
We now want to see that this construction works.
For (i) we know, by the above lemma, that exactly one $p_{[u]}$ holds in any given forcing extension of $\mathcal{M}$. Such an extension satisfies exactly one pattern A of switches and A is an element of exactly one A_w because the A_w form a partition. So exactly one s_w is true for $w \in [u]$. Thus in any forcing extension exactly one $p_w = p_{[u]} \wedge s_w$ holds (for $w \in [u]$).
By construction $A_{w_0} = \{B\}$, the pattern of switches that happens to hold in $\mathcal{M}$, so $\mathcal{M} \models s_{w_0}$ and by the above lemma it follows that $\mathcal{M} \models p_{w_0}$.
To see (ii) suppose $\mathcal{M}[G] \models p_w$ for some filter G. Then $\mathcal{M}[G] \models p_{[w]} \wedge s_w$, so if $w \leq_{\mathcal{F}} v$ we already know that $p_{[v]}$ is forceable over $\mathcal{M}[G]$ by the the above lemma. s_v is always forceable without affecting the buttons (independence!), so $p_{[v]} \wedge s_v$ is forceable over $\mathcal{M}[G]$.
On the other hand, if p_v is forceable over $\mathcal{M}[G]$, then $p_{[v]}$ is forceable, i.e. $[w] \leq_{\mathcal{F}/\equiv} [v]$ which implies $w \leq_{\mathcal{F}} v$. □

We will now prove a lemma that is a little stronger than the lemma mentioned at the beginning of this chapter. A following corollary will then prove this first lemma.

Lemma 3.0.8 *Let $\mathcal{N} = (\mathcal{F}, \mathcal{V})$ be a Kripke Model, $\mathcal{F}$ a finite pre-lattice and $w_0 \in \mathcal{F}$. Let $\mathcal{M}$ be a model of set theory with a sufficiently large independent family of buttons and switches. One can find an assignment $q_i \mapsto \psi_i$ from the propositional variables to set theoretical assertions such that for any modal assertion ϕ:*

$$(\mathcal{N}, w) \models \phi(q_0, ..., q_n) \text{ iff } \mathcal{M} \models \Box(p_w \to \phi(\psi_0, ..., \psi_n))$$

where the p_w are defined as in the above lemma.

Proof. First of all notice that we can trim $\mathcal{F}$ in such a way that we cut off all worlds in $\mathcal{F}$ not accessible from w_0. By doing so we make w_0 a minimal node. Let us therefore assume without loss of generality that w_0 is a minimal node in $\mathcal{F}$. Hence we may assume that $\mathcal{M} \models p_{w_0}$.
Set $\psi_i = \bigvee_{(\mathcal{N},w)\models q_i} p_w$. We now show the claim by induction on the complexity of ϕ.
Let ϕ be atomic. If $\mathcal{M}[G] \models p_w$ then $\mathcal{M}[G] \models \psi_i$ because $(\mathcal{N}, w) \models q_i$ iff p_w is one of the disjuncts of ψ_i. So $\mathcal{M} \models \Box(p_w \to \psi_i)$. Conversely, if $(p_w \to \psi_i)$ is true in some extension where p_w holds, then ψ_i is true in this extension. This implies $(\mathcal{N}, w) \models q_i$.

Suppose the claim holds for ϕ_0 and ϕ_1. Then $(\mathcal{N}, w) \models \phi_0(\vec{q_0}) \wedge \phi_1(\vec{q_1})$ iff $(\mathcal{N}, w) \models \phi_0(\vec{q_0})$ and $(\mathcal{N}, w) \models \phi_1(\vec{q_1})$ iff $\mathcal{M} \models \Box(p_w \rightarrow \phi_0(\vec{\psi_0}))$ and $\mathcal{M} \models \Box(p_w \rightarrow \phi_1(\vec{\psi_1}))$ iff $\mathcal{M} \models \Box(p_w \rightarrow (\phi_0(\vec{\psi_0}) \wedge \phi_1(\vec{\psi_1}))$.
For negation let $(\mathcal{N}, w) \models \neg\phi(q_1, ..., q_n)$. By induction hypothesis we know that $\mathcal{M} \not\models \Box(p_w \rightarrow \phi(\psi_1, ..., \psi_n))$ so we find some partial order $\mathbb{P}$ such that $\mathcal{M}^{\mathbb{P}} \models p_w \wedge \neg\phi(\psi_1, ..., \psi_n)$. But the thruthvalue of the ψ_i necessarily depends only on the truth value of p_v for the various $v \in \mathcal{F}$, so it follows that $p_w \rightarrow \neg\phi(\psi_1, ..., \psi_n)$ holds in every extension.
On the other hand from $\mathcal{M} \models \Box(p_w \rightarrow \neg\phi(\psi_1, ..., \psi_n))$ we derive that $\mathcal{M} \not\models \Box(p_w \rightarrow \psi(\psi_1, ..., \psi_n))$, so by induction hypothesis $(\mathcal{N}, w) \models \neg\phi(q_1, ..., q_n)$.
To see the claim for the modal operator first suppose that $(\mathcal{N}, w) \models \Diamond\phi(q_1, ..., q_n)$. We find a $v \in \mathcal{F}$ with $w \leq_{\mathcal{F}} v$ such that $(\mathcal{N}, v) \models \phi(q_1, ..., q_n)$. By induction $\mathcal{M} \models \Box(p_v \rightarrow \phi(\psi_1, ..., \psi_n))$. But the construction of the p_w's ensured that $\mathcal{M} \models \Box(p_w \rightarrow \Diamond p_v)$ which implies $\mathcal{M} \models \Box(p_w \rightarrow \Diamond\phi(\psi_1, ..., \psi_n))$.
Now suppose $\mathcal{M} \models \Box(p_w \rightarrow \Diamond\phi(\psi_1, ..., \psi_n))$. So $\phi(\psi_1, ..., \psi_n)$ is forceable over any extension of $\mathcal{M}$ satisying p_w. So we find a $v \in \mathcal{F}$ with $w \leq_{\mathcal{F}} v$ and some partial order $\mathbb{P}$ such that $\mathcal{M}^{\mathbb{P}} \models p_v \wedge \phi(\psi_1, ..., \psi_n)$. Again the truth value of the ψ_i depends only on the truth values of the various p_u, so it must be that $\mathcal{M} \models \Box(p_v \rightarrow \phi(\psi_1, ..., \psi_n))$. By induction hypothesis this equals $(\mathcal{N}, v) \models \phi(q_1, ..., q_n)$, so by $w \leq_{\mathcal{F}} v$ we have $(\mathcal{N}, w) \models \Diamond\phi(q_1, ..., q_n)$ as desired. □

Corollary 3.0.9 *With the notation from above: If* $[(\mathcal{N}, w) \models \phi(q_1, ..., q_n)$ *iff* $\mathcal{M} \models \Box(p_w \rightarrow \phi(\psi_1, ..., \psi_n)]$ *then* $[(\mathcal{N}, w_0) \models \phi(q_0, ..., q_n)$ *iff* $\mathcal{M} \models \phi(\psi_1, ..., \psi_n)]$.

Proof. From left to right remember that $\mathcal{M} \models p_{w_0}$ by construction. By assumption $(\mathcal{N}, w_0) \models \phi(q_1, .., q_n)$ implies $\mathcal{M} \models \Box(p_{w_0} \rightarrow \phi(\psi_1, ..., \psi_n))$, so $\mathcal{M} \models \phi(\psi_1, ..., \psi_n)$ because $\mathcal{M}$ is a forcing extension of itself. The converse also holds because if $(\mathcal{N}, w_0) \models \neg\phi(q_1, ..., q_n)$ then the above has shown $\mathcal{M} \models \neg\phi(\psi_1, ..., \psi_n)$. □

Theorem 3.0.10 *The modal logic of forcing is S4.2*

Proof. We have already seen that every assertion in $S4.2$ holds under the forcing interpretation. So let ϕ be a modal assertion not in $S4.2$. Because $S4.2$ is sound and complete for the class of finite pre-lattices we know that there is a Kripke Model $\mathcal{N}$ whose frame $\mathcal{F}$ is a finite pre-lattice and ϕ fails in $\mathcal{N}$ at some world w. Without loss of generality w is an initial world in $\mathcal{N}$.
Assuming consistency of ZFC we know that $ZFC + V = L$ is consistent, so there is a model, namely L, with a sufficiently large family of independent buttons and switches by lemma 3.0.5. So there is an assignment from the propositional variables of ϕ to sentences ψ_i in the language of set theory such that $L \models \neg\phi(\psi_0, ..., \psi_n)$. Therefore ϕ is not a valid principle of forcing in L and hence not a ZFC-provable principle of forcing. □

Chapter 4

Restriction to a class of forcings

In this chapter we will investigate how the modal validities of forcing may change when we restrict ourselves to a class of forcings. First we will have a look at c.c.c.-forcings to point out the difficulties that may occur in these circumstances. Secondly we will consider ω-closed forcings and give prove that the ZFC-provable principles for these forcings equal S4.2.

4.1 Restriction to c.c.c.-forcings

Recall the following definitions.

Definition 4.1.1 A tree T is called *special* if there is a function $f : T \to \omega$ such that $\forall x, y \in T$: $x < y \Rightarrow f(x) \neq f(y)$. A tree T is called *Aronszajn* if $|T| = \omega_1$, $\forall \alpha < \omega_1$ ($|Lev_\alpha T| < \omega_1$) and every chain in T is countable.

Now suppose some $\mathcal{M}$ such that $\mathcal{M} \models$ "T is special" and some partial order $\mathbb{P}$ and a $\mathbb{P}$-generic filter G such that $\mathcal{M}[G] \models$ "T is special and non-Aronszajn". Let $B =< b_\xi \mid \xi < \omega_1^{\mathcal{M}} >$ be a branch through T in $\mathcal{M}[G]$. It follows from "T is special" in $\mathcal{M}[G]$ that there is an f such that for all b_ξ, b_γ with $b_\xi < b_\gamma : f(b_\xi) < f(b_\gamma)$, i.e. $f : B \to \omega$ is injective and hence $(\omega_1^{\mathcal{M}} = |B| = \omega)^{\mathcal{M}[G]}$. So ω_1 was collapsed by $\mathbb{P}$, hence $\mathbb{P}$ did not have the c.c.c.

In L consider a Souslin-tree T. The following lemmata will show that, in any c.c.c-extension of L, it is possible to find c.c.c.-extensions where T is either special or non-Aronszajn. The above shows that no further c.c.c.-forcing makes both statements true at the same time. So ϕ'= "The L-least Souslin-tree is special" and ψ'= "The L-least Souslin-tree is non-Aronszajn" are both forceable and we will see that they are forceably necessary by c.c.c.-forcing. But $\Diamond\Box(\phi' \wedge \psi')$ does not hold in L (as the above has shown).

This is certainly a difference between the ZFC-provable validities of the class of all forcings and the ZFC-provable validities of the class of c.c.c.-forcings as for now we were used to the fact that two forceably necessary statements could be conjuncted by forcing. This was represented by the lattice structure of S4.2, in particular that every two elements of a lattice have a join.
For c.c.c.-forcing the statements ϕ' and ψ' would refer to two worlds with no join, so c.c.c.-forcing does not seem to embody the same structure as the class of all forcings.
This motivates us to ask "What are the ZFC-provable principles of Γ-forcing?" for any class of forcings Γ. There are many open questions here and this thesis is not the right place to answer all of them. However, in the next section we will see a detailed treatment of ω-closed forcings as well as $< \kappa$-closed forcings (for definable κ). Hamkins and Löwe have settled the question for collapse forcing of the form $Col(\omega, \theta)$ for any θ, for which the ZFC-provable validities have proven to be S4.3. The paper is yet to appear, see **??**.

Lemma 4.1.2 *S4.2 implies the directedness axiom, i.e.*

$$S4.2 \vdash \Diamond\Box\phi \wedge \Diamond\Box\psi \to \Diamond\Box(\phi \wedge \psi)$$

Proof. To show: For every S4.2-frame $\mathcal{F}$ and every world w in $\mathcal{F}$ if $\mathcal{F} \models \Diamond\Box\phi \wedge \Diamond\Box\psi$ then $\mathcal{F}, w \models \Diamond\Box(\phi \wedge \psi)$.
By the results of chapter 1 assume that $\mathcal{F}$ is a directed partial order.
$\mathcal{F}, w \models \Diamond\Box\phi \wedge \Diamond\Box\psi$ implies $\mathcal{F}, w \models \Diamond\Box\phi$ and $\mathcal{F}, w \models \Diamond\Box\psi$. So there are u, v such that $w\mathcal{R}u$ and $w\mathcal{R}v$ and $\mathcal{F}, u \models \Box\phi$ and $\mathcal{F}, v \models \Box\psi$. Because $\mathcal{R}$ is directed there is a z such that $u\mathcal{R}z$ and $v\mathcal{R}z$. Hence $\mathcal{F}, z \models \phi$ and $\mathcal{F}, z \models \psi$.
Let $z\mathcal{R}z'$. By transitivity of $\mathcal{R}$ we have $u\mathcal{R}z'$ and $v\mathcal{R}z'$ and hence $\mathcal{F}, z' \models \phi \wedge \psi$. So $\mathcal{F}, z \models \Box(\phi \wedge \psi)$. Using transitivity again we see $w\mathcal{R}z$ and hence $\mathcal{F}, w \models \Diamond\Box(\phi \wedge \psi)$. □

Once we have shown that ϕ' and ψ' are forceably necessary the directedness axiom shows that not all of S4.2 holds for c.c.c.-forcing.

Lemma 4.1.3 *Let T be Souslin in $\mathcal{M}$. There is a partial order which satisfies the c.c.c. such that in $\mathcal{M}^{\mathbb{P}}$ T is no longer Aronszajn.*

Proof. Take $\mathbb{P}$ to be the inverse order on T. Then $\mathbb{P}$ satisfies the c.c.c. because T was Souslin and T is no longer Aronszajn because $\mathbb{P}$ adds a branch through T. □

Thus ψ' is forceable over L.

Lemma 4.1.4 *Let T be Aronszajn in $\mathcal{M}$. Then there is a partial order $\mathbb{P}$ which satisfies the c.c.c. such that T is special in the $\mathcal{M}[G]$ for every G $\mathbb{P}$-generic over $\mathcal{M}$.*

Proof. Define $\mathbb{P}$ as a set of functions such that for all $p \in \mathbb{P}$:

(i) $dom(p) \subset_{fin} T$

(ii) $ran(p) \subset \omega$

(iii) if $x, y \in dom(p) \wedge x \parallel y \Rightarrow p(x) \neq p(y)$

and order $\mathbb{P}$ by end-extension. Then for each $x \in T$ the set

$$D_x = \{p \in \mathbb{P} \mid x \in dom(p)\}$$

is clearly dense in $\mathbb{P}$, so for any $\mathbb{P}$-generic filter G the domain of $f = \bigcup G$ equals T, so f maps T into ω. Now let $x, y \in f^{-1}(n)$ for any $n < \omega$ and $x \parallel y$ in T. By *(iii)*: $p(x) \neq p(y)$ so $f(x) \neq f(y)$, a contradiction. Hence every f^{-1} forms an antichain. So f witnesses that T is special (see [1]).

It only remains to show that $\mathbb{P}$ has the c.c.c.. We need the following claim
Claim 1. If T is an Aronszajn tree and W an uncountable collection of finite pairwise disjoint subsets of T, then there exists $S, S' \in W$ such that any $x \in S$ is incomparable with any $y \in S'$.
Proof of Claim 1. Without loss of generality assume that, for some $n < \omega$, $|S| = n$ for all $S \in W$ (because $|S| = n$ holds for uncountably many $S \in W$).
Fix an enumeration $\{z_1, ..., z_n\}$ for each $S \in W$. Because W is uncountable there are ultrafilters on W whose elements are uncountable (see [2]). Let D be such an ultrafilter.
Assume the claim is false. For each $x \in T$ and each $k \in \{1, ..., n\}$ let $Y_{x,k}$ be the set of all $S \in W$ such that x is comparable with the k-th element of S. By assumption, any S, S' contains comparable elements, so

$$\bigcup_{x \in S} \bigcup_{k=1}^{n} Y_{x,k} = W$$

holds for every $S \in W$. Because S is finite at least one of the $Y_{x,k}$ has to be uncountable. Because D is an ultrafilter we may pick, for any $S \in W$, an $x = x_S$, an element of S, and a $k = k_S$ such that $Y_{x,k} \in D$. Now define

$$Z_k = \{S \in W \mid k_S = k\} \text{ for } k \leq n$$

Then there is a k' such that $Z_{k'}$ is uncountable. This will lead to a contradiction because we will show that the elements of $\{x_S \mid S \in Z_{k'}\}$ are pairwise comparable i.e. generate an uncountable branch through T.
Let $S_1, S_2 \in Z_{k'}$, $x = x_{S_1}$, $y = x_{S_2}$. Then $Y_{x,k'}$ and $Y_{y,k'}$ are elements of D, so $Y = Y_{x,k'} \cap Y_{y,k'} \in D$. Therefore Y is uncountable. So for any $S \in Y$ the k'-th element of S is comparable with x and y. x and y are elements of an Aronszajn tree so they have at most countably many predecessors. But Y is uncountable, so there is some $S \in Y$ such that the k'-th element of S is greater than both x and y. But then the k'-th element of S witnesses that x and y are comparable.

[1] $T = \bigcup_{n<\omega} f^{-1}(n)$, i.e. T is a countable union of antichains. One can now define a function g that sends every $x \in T$ to the minimal n such that $x \in f^{-1}(n)$ and $\forall y <_T x$ with $g(y) \in f^{-1}(m)$: $m < n$.

[2] This is a result of Pospíšil, see [7], Theorem 7.6

$\Box$(Claim 1)

Now let W be an uncountable subset of $\mathbb{P}$. Then $\{dom(p) \mid p \in W\}$ is uncountable because if not there would be uncountably many functions from a finite set into ω. Using the Δ-system-lemma we find an uncountable $W_1 \subset W$ and a finite $S \subset T$ such that $dom(p) \cap dom(q) = S$ for all $p, q \in W_1$ distinct. We then find a $W_2 \subset W_1$ uncountable such that $p \restriction S = q \restriction S$ for all $p, q \in W_2$ by pigeonhole. So by the above claim there are $p, q \in W_2$ such that any $x \in dom(p) \setminus S$ is incomparable with any $y \in dom(q) \setminus S$. Then $p \cup q \in \mathbb{P}$, i.e. $p \parallel q$ in $\mathbb{P}$, hence $\mathbb{P}$ has the c.c.c.

$\Box$

So we have proven that ϕ' and ψ' are both forceable over L. To see that once ϕ' or ψ' holds it stays true in all further c.c.c. extensions note that a tree that is Aronszajn in the ground model and was made non-Aronszajn in a c.c.c. extension has an uncountable branch. To make the tree Aronszajn again one would have to make the branch countable again, i.e. collapse ω_1. This can not be done by c.c.c. forcing. So ψ' is forceably necessary by c.c.c.-forcing.
ϕ' stays true in all further c.c.c. extension because the f that witnessed that T was special is definable and hence an element of every further forcing extension.

Let us write $\Diamond_\Gamma \phi$ and $\Box_\Gamma \phi$ when we mean that ϕ is forceable or necessary by Γ-forcing respectively. So $\Diamond_{c.c.c.}\Box_{c.c.c.}\phi'$ and $\Diamond_{c.c.c.}\Box_{c.c.c.}\psi'$ hold. Denote the ZFC-provable principles of Γ-focing by $Force_\Gamma$. Thus far we have shown in this chapter

$$S4.2 \nsubseteq Force_{c.c.c.}$$

because the directedness axiom does not hold.

Let us quickly show that:

Lemma 4.1.5 *S4* $\subseteq Force_{c.c.c.}$

Proof. To see (K) suppose ϕ and $\phi \rightarrow \psi$ holds in all c.c.c.-extensions. Then ψ holds in all c.c.c. extensions, so $\Box_{c.c.c.}(\phi \rightarrow \psi) \rightarrow (\Box_{c.c.c.}\phi \rightarrow \Box_{c.c.c.}\psi)$ holds.
For (Dual) note that if there is no c.c.c. extension such that ϕ holds, then in every c.c.c.-extension $\neg\phi$. On the other hand, if in every c.c.c. extension $\neg\phi$ holds, then there is no c.c.c.-extension where ϕ holds. Hence $\neg\Diamond_{c.c.c.}\phi \leftrightarrow \Box_{c.c.c.}\neg\phi$.
For (S) we use the same argument as for the class of all forcings: Every model of ZFC is a (c.c.c.)-extension of itself, so $\Box_{c.c.c.}\phi \rightarrow \phi$.
(4) holds because if $\mathbb{P}$ has the c.c.c. in the ground model $\mathcal{M}$ and in $\mathcal{M}[G]$, where G is a $\mathbb{P}$-generic filter over $\mathcal{M}$, there is a partial order $\mathbb{Q}$ that satisfies the c.c.c., then, because ω_1 is a regular uncountable cardinal, $\mathbb{P} * \dot{\mathbb{Q}}$ satisfies the c.c.c. in $\mathcal{M}$ (see [7] Theorem 16.4 for details). Hence $\Box_{c.c.c.}\phi \rightarrow \Box_{c.c.c.}\Box_{c.c.c.}\phi$ holds.
$Force_{c.c.c.}$ is clearly closed under modus ponens and uniform substitution and if ϕ holds in every model of set theory, then so does $\Box_{c.c.c.}\phi$. $\Box$

4.2 The modal logic of ω-closed forcing

As seen in the previous section, restricting oneself to a class of forcings may result in various difficulties, for the internal structure of forcing may change drastically. However, for ω-closed forcing we manage to apply the same algorithm used to show $Force = S4.2$ fruitfully. This is the main result of this thesis.

Theorem 4.2.1 (Main Theorem) $Force_{\omega-closed} = S4.2$

Proof. To see the right to left inclusion one argues as we have done for arbitrary forcing. Note that, unlike c.c.c.-forcing, the product of two ω-closed forcings is again ω-closed, hence the proof for (.2) works as well.

For the left to right inclusion note that every switch we defined was especially ω-closed. There was a remark on this when we saw the independence of the switches.

For all $n > 0$ the buttons ϕ_n defined in section 2.3 are ω-closed buttons because we pushed a button by forcing with $Col(\omega_{3n}^L, \omega_{3n+1}^L)$. This forcing is $< \omega_{3n}^L$-closed and therefore, for any set $A \subset_{fin} (\omega \setminus \{0\})$, the family $\{\phi_n \mid n \in A\}$ forms an arbitrarily large but finite family of ω-closed buttons in L. Note that, if $n > 0$, the case $(|\omega_{3n}^L| = \omega)^{\mathcal{M}}$ were $\mathcal{M}$ is some ω-closed extension of L may not occur because ω-closed forcing can not collaps an uncountable cardinal onto ω.

Hence the buttons and switches are ω-closed and section 2 has already shown their independence. Therefore we argue just as in 3.0.10 to conclude the theorem.

□

4.3 The modal logic of $< \kappa$-closed forcing

In this section we want to examine what happens when we extend our class of forcings to $< \kappa$-closed forcings.

A first and obviously important question is: "What is meant by κ?". We know that κ may well change its cardinality by forcing and the property to be $< \delta$-closed was only defined for cardinals δ. Therefore it does not seem to make sense to talk about the modal logic of $< \kappa$-closed forcings, unless we specify what is meant by κ.

The easiest way to deal with this difficulty is to request κ to be definable by some formula ϕ. As an example one may think of a ϕ that says "x is the third uncountable cardinal". One then calculates κ in the various extensions using ϕ and looks at the class of $< \kappa$-closed forcings calculated in the model. Note that we did this already for ω-closed forcing, only that we did not mention the formula ϕ that defines ω because it is clear that such a formula exists.

Our goal is to show the follwing theorem.

Theorem 4.3.1 *Let κ be a cardinal definable by some formula ϕ. Then the modal logic of $< \kappa$-closed forcing is S4.2*

Obviously $S4.2 \subseteq Force_{<\kappa-closed}$ for one uses exactly the same arguments presented to show this for ω-closed forcing. Therefore only the other inclusion remains to be shown. As we already know by now this boils down to the task of finding independent families of buttons and switches. These families will be basically the same as in the previous section with only minor changes to the buttons and switches. The idea is to arrange the buttons and switches "high enough" to make them $< \kappa$-closed switches while keeping them independent from one another. More precisely define

$$\phi_n \equiv \aleph^L_{\kappa+3n} \notin Card \wedge \aleph^L_{\kappa+3n+1} \notin Card \wedge \aleph^L_{\kappa+3n+2} \notin Card \wedge \left|\mathscr{P}(\aleph^L_{\kappa+3n})\right| > \aleph^L_{\kappa+3n+1}$$

and

$$\psi_m \equiv 2^{\aleph_{\aleph_{\kappa+2m+1}}} = (\aleph_{\aleph_{\kappa+2m+1}})^+$$

We claim that for any $A, B \subset_{fin} \omega$ the set $\{\phi_n \mid n \in A\} \cup \{\psi_m \mid m \in B\}$ is an independent family of $< \kappa$-closed buttons and switches in L.

Let us consider the buttons first. We have to see that they are in fact $< \kappa$-closed buttons and that they are independent.

Suppose some $< \kappa$-closed extension $\mathcal{M}$ of L where $\Box\phi_n$ does not hold. By precisely the same argument used in 2.3 we see that $\mathcal{M} \models \neg\phi_n$. In order to push ϕ_n we force with $Col(\aleph^L_{\kappa+3n}, \aleph^L_{\kappa+3n+1})$. Because $\neg\phi_n$ holds in $\mathcal{M}$ we have that $\aleph^L_{\kappa+3n}$ is still a cardinal in $\mathcal{M}$. It follows that the forcing is $< \aleph^L_{\kappa+3n}$-closed in $\mathcal{M}$. Notice that $\mathcal{M}$ is a $< \kappa$-closed forcing extension of L, so L and $\mathcal{M}$ have the same cardinals below κ. We then see that $(\kappa \leq \left|\aleph^L_{\kappa+3n}\right|)^{\mathcal{M}}$, so the forcing defined in $\mathcal{M}$ is especially $< \kappa$-closed. Therefore the buttons are $< \kappa$-closed buttons.

To show their independence first notice that no ϕ_n holds in L by $L \models GCH$. To see that the ϕ_n can be pushed independently we may now use the proof presented in 2.4 where we replace any occurance of ω^L_i by $\aleph^L_{\kappa+i}$. Independence follows easily.

Now consider the switches. To turn ψ_n on we force with $Fn((\aleph_{\aleph_{\kappa+2n+1}})^{++} \times \aleph_{\aleph_{\kappa+2n+1}}, 2, \aleph_{\aleph_{\kappa+2n+1}})$. By the results of chapter 2 this forcing is $< \aleph_{\aleph_{\kappa+2n+1}}$-closed (note especially that $\aleph_{\aleph_{\kappa+2n+1}}$ is a successor and hence regular).

To turn ψ_n off we use collapse-forcing of the form $Col(\aleph_{\aleph_{\kappa+2n+1}}, \overline{\overline{\mathscr{P}(\aleph_{\aleph_{\kappa+2n+1}})}})$. This is $< \aleph_{\aleph_{\kappa+2n+1}}$-closed as well. By $(\kappa \leq \aleph_{\aleph_{\kappa+2n+1}})^{\mathcal{M}}$, for any $< \kappa$-closed forcing extension $\mathcal{M}$ of L, this has shown that ψ_n is a $< \kappa$-closed switch.

Modulo minor changes the independence of the switches was proven by the general results in chapter 2.

Independence of the whole family uses the same argument as independence for the ω-closed family:
The switches are $< \aleph_{\aleph_\kappa}$-closed, hence no new subsets are added below $\aleph_{\aleph_\kappa}$, which is to say: The switches do not affect the truth value of the buttons.

The forcings used to turn on the buttons on the other hand have the $\aleph_{\aleph_\kappa}$-c.c. so "not many" subsets are added to cardinals above $\aleph_{\aleph_\kappa}$, i.e. the GCH above $\aleph_{\aleph_\kappa}$ does not get affected.

We have hence shown that there is an independent family of $< \kappa$-closed buttons and switches in L. This in turn proves the above stated theorem.

Bibliography

[1] Patrick Blackburn; Maarten de Rijke; and Yde Venema, *Modal Logic*, Cambridge University Press, 2001

[2] Alexander Chagrov; Michael Zakharyaschev, *Modal logic*, Claredon Press, Oxford 1997

[3] Keith Devlin, *Constructibility*, Perspectives in Mathematical Logic, Springer Verlag, 1984

[4] Joel David Hamkins, *A simple maximality principle*, Journal of Symbolic Logic, June 2003

[5] Joel David Hamkins, Benedikt Löwe, 2008. *The Modal Logic of Forcing*, Transactions of the American Mathematical Society, volume 360, number 4

[6] Joel David Hamkins, Benedikt Löwe, *Structural relations between classes of forcings and their modal logics*, in preparation

[7] Thomas Jech, *Set Theory*, Third edition, Springer Verlag 2006

[8] Kenneth Kunen, *Set Theory, An Introduction to Independence Proofs*, Studies in logic and the foundations of mathematics, Elsevier, 2006

[9] Graham Priest, *An introduction to non-classical logic*, Cambridge University Press, 2001

[10] Ralf Schindler, *Logic I-IV Skript*, available at wwwmath1.uni-muenster.de/logik/Personen/rds/