Chapter udf

Filtrations and Decidability

fil.1 Introduction

nml:fil:int:sec One important question about a logic is always whether it is decidable, i.e., if there is an effective procedure which will answer the question "is this formula valid." Propositional logic is decidable: we can effectively test if a formula is a tautology by constructing a truth table, and for a given formula, the truth table is finite. But we can't obviously test if a modal formula is true in all models, for there are infinitely many of them. We can list all the finite models relevant to a given formula, since only the assignment of subsets of worlds to propositional variables which actually occur in the formula are relevant. If the accessibility relation is fixed, the possible different assignments $V(p)$ are just all the subsets of W, and if $|W| = n$ there are 2^n of those. If our formula φ contains m propositional variables there are then 2^{nm} different models with n worlds. For each one, we can test if φ is true at all worlds, simply by computing the truth value of φ in each. Of course, we also have to check all possible accessibility relations, but there are only finitely many relations on n worlds as well (specifically, the number of subsets of $W \times W$, i.e., 2^{n^2}.

If we are not interested in the logic **K**, but a logic defined by some class of models (e.g., the reflexive transitive models), we also have to be able to test if the accessibility relation is of the right kind. We can do that whenever the frames we are interested in are definable by modal formulas (e.g., by testing if T and 4 valid in the frame). So, the idea would be to run through all the finite frames, test each one if it is a frame in the class we're interested in, then list all the possible models on that frame and test if φ is true in each. If not, stop: φ is not valid in the class of models of interest.

There is a problem with this idea: we don't know when, if ever, we can stop looking. If the formula has a finite countermodel, our procedure will find it. But if it has no finite countermodel, we won't get an answer. The formula may be valid (no countermodels at all), or it have only an infinite countermodel, which we'll never look at. This problem can be overcome if we can show that every formula that has a countermodel has a finite countermodel. If this is the case we say the logic has the *finite model property*.

But how would we show that a logic has the finite model property? One way of doing this would be to find a way to turn an infinite (counter)model of φ into a finite one. If that can be done, then whenever there is a model in which φ is not true, then the resulting finite model also makes φ not true. That finite model will show up on our list of all finite models, and we will eventually determine, for every formula that is not valid, that it isn't. Our procedure won't terminate if the formula is valid. If we can show in addition that there is some maximum size that the finite model our procedure provides can have, and that this maximum size depends only on the formula φ, we will have a size up to which we have to test finite models in our search for countermodels. If we haven't found a countermodel by then, there are none. Then our procedure will, in fact, decide the question "is φ valid?" for any formula φ.

A strategy that often works for turning infinite structures into finite structures is that of "identifying" elements of the structure which behave the same way in relevant respects. If there are infinitely many worlds in $\mathfrak{M}$ that behave the same in relevant respects, then we might hope that there are only finitely many "classes" of such worlds. In other words, we partition the set of worlds in the right way. Each partition contains infinitely many worlds, but there are only finitely many partitions. Then we define a new model $\mathfrak{M}^*$ where the worlds are the partitions. Finitely many partitions in the old model give us finitely many worlds in the new model, i.e., a finite model. Let's call the partition a world w is in $[w]$. We'll want it to be the case that $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \varphi$, since we want the new model to be a countermodel to φ if the old one was. This requires that we define the partition, as well as the accessibility relation of $\mathfrak{M}^*$ in the right way.

To see how this would go, first imagine we have no accessibility relation. $\mathfrak{M}, w \Vdash \Box\psi$ iff for some $v \in W$, $\mathfrak{M}, v \Vdash \Box\psi$, and the same for $\mathfrak{M}^*$, except with $[w]$ and $[v]$. As a first idea, let's say that two worlds u and v are equivalent (belong to the same partition) if they agree on all propositional variables in $\mathfrak{M}$, i.e., $\mathfrak{M}, u \Vdash p$ iff $\mathfrak{M}, v \Vdash p$. Let $V^*(p) = \{[w] : \mathfrak{M}, w \Vdash p\}$. Our aim is to show that $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \varphi$. Obviously, we'd prove this by induction: The base case would be $\varphi \equiv p$. First suppose $\mathfrak{M}, w \Vdash p$. Then $[w] \in V^*$ by definition, so $\mathfrak{M}^*, [w] \Vdash p$. Now suppose that $\mathfrak{M}^*, [w] \Vdash p$. That means that $[w] \in V^*(p)$, i.e., for some v equivalent to w, $\mathfrak{M}, v \Vdash p$. But "w equivalent to v" means "w and v make all the same propositional variables true," so $\mathfrak{M}, w \Vdash p$. Now for the inductive step, e.g., $\varphi \equiv \neg\psi$. Then $\mathfrak{M}, w \Vdash \neg\psi$ iff $\mathfrak{M}, w \nVdash \psi$ iff $\mathfrak{M}^*, [w] \nVdash \psi$ (by inductive hypothesis) iff $\mathfrak{M}^*, [w] \Vdash \neg\psi$. Similarly for the other non-modal operators. It also works for $\Box$: suppose $\mathfrak{M}^*, [w] \Vdash \Box\psi$. That means that for every $[u]$, $\mathfrak{M}^*, [u] \Vdash \psi$. By inductive hypothesis, for every u, $\mathfrak{M}, u \Vdash \psi$. Consequently, $\mathfrak{M}, w \Vdash \Box\psi$.

In the general case, where we have to also define the accessibility relation for $\mathfrak{M}^*$, things are more complicated. We'll call a model $\mathfrak{M}^*$ a *filtration* if its accessibility relation R^* satisfies the conditions required to make the inductive proof above go through. Then any filtration $\mathfrak{M}^*$ will make φ true at $[w]$ iff $\mathfrak{M}$ makes φ true at w. However, now we also have to show that there *are* filtrations, i.e., we can define R^* so that it satisfies the required conditions. In order for this to work, however, we have to require that worlds u, v count as equivalent not just when they agree on all propositional variables, but on all sub-formulas of φ. Since φ has only finitely many sub-formulas, this will still guarantee that the filtration is finite. There is not just one way to define a filtration, and in order to make sure that the accessibility relation of the filtration satisfies the required properties (e.g., reflexive, transitive, etc.) we have to be inventive with the definition of R^*.

fil.2 Preliminaries

nml:fil:pre: sec

Filtrations allow us to establish the decidability of our systems of modal logic by showing that they have the *finite model property*, i.e., that any formula that is true (false) in a model is also true (false) in a *finite* model. Filtrations are defined relative to sets of formulas which are closed under subformulas.

nml:fil:pre: defn:modallyclosed

Definition fil.1. A set Γ of formulas is *closed under subformulas* if it contains every subformula of a formula in Γ. Further, Γ is *modally closed* if it is closed under subformulas and moreover $\varphi \in \Gamma$ implies $\Box\varphi, \Diamond\varphi \in \Gamma$.

For instance, given a formula φ, the set of all its sub-formulas is closed under sub-formulas. When we're defining a filtration of a model through the set of sub-formulas of φ, it will have the property we're after: it makes φ true (false) iff the original model does.

The set of worlds of a filtration of $\mathfrak{M}$ through Γ is defined as the set of all equivalence classes of the following equivalence relation.

Definition fil.2. Let $\mathfrak{M} = \langle W, R, V\rangle$ and suppose Γ is closed under sub-formulas. Define a relation $\equiv$ on W to hold of any two worlds that make the same formulas from Γ true, i.e.:

$$u \equiv v \quad \text{if and only if} \quad \forall\varphi \in \Gamma : \mathfrak{M}, u \Vdash \varphi \Leftrightarrow \mathfrak{M}, v \Vdash \varphi.$$

The equivalence class $[w]_\equiv$ of a world w, or $[w]$ for short, is the set of all worlds $\equiv$-equivalent to w:

$$[w] = \{v : v \equiv w\}.$$

Proposition fil.3. *Given $\mathfrak{M}$ and Γ, $\equiv$ as defined above is an equivalence relation, i.e., it is reflexive, symmetric, and transitive.*

Proof. The relation $\equiv$ is reflexive, since w makes exactly the same formulas from Γ true as itself. It is symmetric since if u makes the same formulas from Γ true as v, the same holds for v and u. It is also transitive, since if u makes the same formulas from Γ true as v, and v as w, then u makes the same formulas from Γ true as w. □

The relation $\equiv$, like any equivalence relation, divides W into *partitions*, i.e., subsets of W which are pairwise disjoint, and together cover all of W. Every $w \in W$ is an element of one of the partitions, namely of $[w]$, since $w \equiv w$. So the partitions $[w]$ cover all of W. They are pairwise disjoint, for if $u \in [w]$ and $u \in [v]$, then $u \equiv w$ and $u \equiv v$, and by symmetry and transitivity, $w \equiv v$, and so $[w] = [v]$.

fil.3 Filtrations

Rather than define "the" filtration of $\mathfrak{M}$ through Γ, we define when a model $\mathfrak{M}^*$ counts as a filtration of $\mathfrak{M}$. All filtrations have the same set of worlds W^* and the same valuation V^*. But different filtrations may have different accessibility relations R^*. To count as a filtration, R^* has to satisfy a number of conditions, however. These conditions are exactly what we'll require to prove the main result, namely that $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \varphi$, provided $\varphi \in \Gamma$. nml:fil:fil:sec

Definition fil.4. Let Γ be closed under subformulas and $\mathfrak{M} = \langle W, R, V \rangle$. A *filtration of* $\mathfrak{M}$ *through* Γ is any model $\mathfrak{M}^* = \langle W^*, R^*, V^* \rangle$, where: nml:fil:fil:defn:filtration

1. $W^* = \{[w] : w \in W\}$;
2. For any $u, v \in W$: nml:fil:fil:defn:filtration-R
    a) If Ruv then $R^*[u][v]$; nml:fil:fil:defn:filtration-R1
    b) If $R^*[u][v]$ then for any $\Box\varphi \in \Gamma$, if $\mathfrak{M}, u \Vdash \Box\varphi$ then $\mathfrak{M}, v \Vdash \varphi$; nml:fil:fil:defn:filtration-R2
    c) If $R^*[u][v]$ then for any $\Diamond\varphi \in \Gamma$, if $\mathfrak{M}, v \Vdash \varphi$ then $\mathfrak{M}, u \Vdash \Diamond\varphi$. nml:fil:fil:defn:filtration-R3
3. $V^*(p) = \{[u] : u \in V(p)\}$.

It's worthwhile thinking about what $V^*(p)$ is: the set consisting of the equivalence classes $[w]$ of all worlds w where p is true in $\mathfrak{M}$. On the one hand, if $w \in V(p)$, then $[w] \in V^*(p)$ by that definition. However, it is not necessarily the case that if $[w] \in V^*(p)$, then $w \in V(p)$. If $[w] \in V^*(p)$ we are only guaranteed that $[w] = [u]$ for *some* $u \in V(p)$. Of course, $[w] = [u]$ means that $w \equiv u$. So, when $[w] \in V^*(p)$ we can (only) conclude that $w \equiv u$ for some $u \in V(p)$.

Theorem fil.5. *If $\mathfrak{M}^*$ is a filtration of $\mathfrak{M}$ through Γ, then for every $\varphi \in \Gamma$ and $w \in W$, we have $\mathfrak{M}, w \Vdash \varphi$ if and only if $\mathfrak{M}^*, [w] \Vdash \varphi$.* *nml:fil:fil:thm:filtrations*

Proof. By induction on φ, using the fact that Γ is closed under subformulas. Since $\varphi \in \Gamma$ and Γ is closed under sub-formulas, all sub-formulas of φ are also $\in \Gamma$. Hence in each inductive step, the induction hypothesis applies to the sub-formulas of φ.

1. $\varphi \equiv \bot$: Neither $\mathfrak{M}, w \Vdash \varphi$ nor $\mathfrak{M}^*, w \Vdash \varphi$.
2. $\varphi \equiv \top$: Both $\mathfrak{M}, w \Vdash \varphi$ and $\mathfrak{M}^*, w \Vdash \varphi$.

3. $\varphi \equiv p$: The left-to-right direction is immediate, as $\mathfrak{M}, w \Vdash \varphi$ only if $w \in V(p)$, which implies $[w] \in V^*(p)$, i.e., $\mathfrak{M}^*, [w] \Vdash \varphi$. Conversely, suppose $\mathfrak{M}^*, [w] \Vdash \varphi$, i.e., $[w] \in V^*(p)$. Then for some $v \in V(p)$, $w \equiv v$. Of course then also $\mathfrak{M}, v \Vdash p$. Since $w \equiv v$, w and v make the same formulas from Γ true. Since by assumption $p \in \Gamma$ and $\mathfrak{M}, v \Vdash p$, $\mathfrak{M}, w \Vdash \varphi$.

4. $\varphi \equiv \neg\psi$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \nVdash \psi$. By induction hypothesis, $\mathfrak{M}, w \nVdash \psi$ iff $\mathfrak{M}^*, [w] \nVdash \psi$. Finally, $\mathfrak{M}^*, [w] \nVdash \psi$ iff $\mathfrak{M}^*, [w] \Vdash \varphi$.

5. $\varphi \equiv (\psi \land \chi)$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \Vdash \psi$ and $\mathfrak{M}, w \Vdash \chi$. By induction hypothesis, $\mathfrak{M}, w \Vdash \psi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$, and $\mathfrak{M}, w \Vdash \chi$ iff $\mathfrak{M}^*, [w] \Vdash \chi$. And $\mathfrak{M}^*, [w] \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$ and $\mathfrak{M}^*, [w] \Vdash \chi$.

6. $\varphi \equiv (\psi \lor \chi)$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \Vdash \psi$ or $\mathfrak{M}, w \Vdash \chi$. By induction hypothesis, $\mathfrak{M}, w \Vdash \psi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$, and $\mathfrak{M}, w \Vdash \chi$ iff $\mathfrak{M}^*, [w] \Vdash \chi$. And $\mathfrak{M}^*, [w] \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$ or $\mathfrak{M}^*, [w] \Vdash \chi$.

7. $\varphi \equiv (\psi \to \chi)$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \nVdash \psi$ or $\mathfrak{M}, w \Vdash \chi$. By induction hypothesis, $\mathfrak{M}, w \Vdash \psi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$, and $\mathfrak{M}, w \Vdash \chi$ iff $\mathfrak{M}^*, [w] \Vdash \chi$. And $\mathfrak{M}^*, [w] \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \nVdash \psi$ or $\mathfrak{M}^*, [w] \Vdash \chi$.

8. $\varphi \equiv (\psi \leftrightarrow \chi)$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \Vdash \psi$ and $\mathfrak{M}, w \Vdash \chi$, or $\mathfrak{M}, w \nVdash \psi$ and $\mathfrak{M}, w \nVdash \chi$. By induction hypothesis, $\mathfrak{M}, w \Vdash \psi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$, and $\mathfrak{M}, w \Vdash \chi$ iff $\mathfrak{M}^*, [w] \Vdash \chi$. And $\mathfrak{M}^*, [w] \Vdash \varphi$ iff $\mathfrak{M}^*, [w] \Vdash \psi$ and $\mathfrak{M}^*, [w] \Vdash \chi$, or $\mathfrak{M}^*, [w] \nVdash \psi$ and $\mathfrak{M}^*, [w] \nVdash \chi$.

9. $\varphi \equiv \Box\psi$: Suppose $\mathfrak{M}, w \Vdash \varphi$; to show that $\mathfrak{M}^*, [w] \Vdash \varphi$, let v be such that $R^*[w][v]$. From Definition fil.4(2b), we have that $\mathfrak{M}, v \Vdash \psi$, and by inductive hypothesis $\mathfrak{M}^*, [v] \Vdash \psi$. Since v was arbitrary, $\mathfrak{M}^*, [w] \Vdash \varphi$ follows.

   Conversely, suppose $\mathfrak{M}^*, [w] \Vdash \varphi$ and let v be arbitrary such that Rwv. From Definition fil.4(2a), we have $R^*[w][v]$, so that $\mathfrak{M}^*, [v] \Vdash \psi$; by inductive hypothesis $\mathfrak{M}, v \Vdash \psi$, and since v was arbitrary, $\mathfrak{M}, u \Vdash \varphi$.

10. $\varphi \equiv \Diamond\psi$: Suppose $\mathfrak{M}, w \Vdash \varphi$. Then for some $v \in W$, Rwv and $\mathfrak{M}, v \Vdash \psi$. By inductive hypothesis $\mathfrak{M}^*, [v] \Vdash \psi$, and by Definition fil.4(2a), we have $R^*[w][v]$. Thus, $\mathfrak{M}^*, [w] \Vdash \varphi$.

    Now suppose $\mathfrak{M}^*, [w] \Vdash \varphi$. Then for some $[v] \in W^*$ with $R^*[w][v]$, $\mathfrak{M}^*, [v] \Vdash \psi$. By inductive hypothesis $\mathfrak{M}, v \Vdash \psi$. By Definition fil.4(2c), we have that $\mathfrak{M}, w \Vdash \varphi$. □

Problem fil.1. Complete the proof of Theorem fil.5

What holds for truth at worlds in a model also holds for truth in a model and validity in a class of models.

Corollary fil.6. *Let Γ be closed under subformulas. Then:*

1. *If $\mathfrak{M}^*$ is a filtration of $\mathfrak{M}$ through Γ then for any $\varphi \in \Gamma$: $\mathfrak{M} \Vdash \varphi$ if and only if $\mathfrak{M}^* \Vdash \varphi$.*

2. *If $\mathcal{C}$ is a class of models and $\Gamma(\mathcal{C})$ is the class of Γ-filtrations of models in $\mathcal{C}$, then any formula $\varphi \in \Gamma$ is valid in $\mathcal{C}$ if and only if it is valid in $\Gamma(\mathcal{C})$.*

fil.4 Examples of Filtrations

We have not yet shown that there are any filtrations. But indeed, for any model $\mathfrak{M}$, there are many filtrations of $\mathfrak{M}$ through Γ. We identify two, in particular: the finest and coarsest filtrations. Filtrations of the same models will differ in their accessibility relation (as Definition fil.4 stipulates directly what W^* and V^* should be). The finest filtration will have as few related worlds as possible, whereas the coarsest will have as many as possible. nml:fil:exf:sec

Definition fil.7. Where Γ is closed under subformulas, the *finest* filtration $\mathfrak{M}^*$ of a model $\mathfrak{M}$ is defined by putting:

$$R^*[u][v] \quad \text{if and only if} \quad \exists u' \in [u]\, \exists v' \in [v] : Ru'v'.$$

Proposition fil.8. *The finest filtration $\mathfrak{M}^*$ is indeed a filtration.* *nml:fil:exf:prop:finest*

Proof. We need to check that R^*, so defined, satisfies Definition fil.4(2). We check the three conditions in turn.

If Ruv then since $u \in [u]$ and $v \in [v]$, also $R^*[u][v]$, so (2a) is satisfied.

For (2b), suppose $\Box\varphi \in \Gamma$, $R^*[u][v]$, and $\mathfrak{M}, u \Vdash \Box\varphi$. By definition of R^*, there are $u' \equiv u$ and $v' \equiv v$ such that $Ru'v'$. Since u and u' agree on Γ, also $\mathfrak{M}, u' \Vdash \Box\varphi$, so that $\mathfrak{M}, v' \Vdash \varphi$. By closure of Γ under sub-formulas, v and v' agree on φ, so $\mathfrak{M}, v \Vdash \varphi$, as desired.

To verify (2c), suppose $\Diamond\varphi \in \Gamma$, $R^*[u][v]$, and $\mathfrak{M}, v \Vdash \varphi$. By definition of R^*, there are $u' \equiv u$ and $v' \equiv v$ such that $Ru'v'$. Since v and v' agree on Γ, and Γ is closed under sub-formulas, also $\mathfrak{M}, v' \Vdash \varphi$, so that $\mathfrak{M}, u' \Vdash \Diamond\varphi$. Since u and u' also agree on Γ, $\mathfrak{M}, u \Vdash \Diamond\varphi$. □

Problem fil.2. Complete the proof of Proposition fil.8.

Definition fil.9. Where Γ is closed under subformulas, the *coarsest* filtration $\mathfrak{M}^*$ of a model $\mathfrak{M}$ is defined by putting $R^*[u][v]$ if and only if *both* of the following conditions are met:

1. If $\Box\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Box\varphi$ then $\mathfrak{M}, v \Vdash \varphi$; nml:fil:exf:defn:coarsest-Box

2. If $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \varphi$ then $\mathfrak{M}, u \Vdash \Diamond\varphi$. nml:fil:exf:defn:coarsest-Diamond

Proposition fil.10. *The coarsest filtration $\mathfrak{M}^*$ is indeed a filtration.*

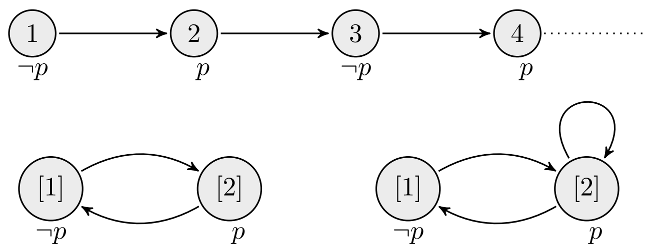

Figure fil.1: An infinite model and its filtrations.

nml:fil:exf:
fig:ex-filtration

Proof. Given the definition of R^*, the only condition that is left to verify is the implication from Ruv to $R^*[u][v]$. So assume Ruv. Suppose $\Box\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Box\varphi$; then obviously $\mathfrak{M}, v \Vdash \varphi$, and (1) is satisfied. Suppose $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \varphi$. Then $\mathfrak{M}, u \Vdash \Diamond\varphi$ since Ruv, and (2) is satisfied. $\Box$

Example fil.11. Let $W = \mathbb{Z}^+$, Rnm iff $m = n + 1$, and $V(p) = \{2n : n \in \mathbb{N}\}$. The model $\mathfrak{M} = \langle W, R, V\rangle$ is depicted in Figure fil.1. The worlds are 1, 2, etc.; each world can access exactly one other world—its successor—and p is true at all and only the even numbers.

Now let Γ be the set of sub-formulas of $\Box p \to p$, i.e., $\{p, \Box p, \Box p \to p\}$. p is true at all and only the even numbers, $\Box p$ is true at all and only the odd numbers, so $\Box p \to p$ is true at all and only the even numbers. In other words, every odd number makes $\Box p$ true and p and $\Box p \to p$ false; every even number makes p and $\Box p \to p$ true, but $\Box p$ false. So $W^* = \{[1], [2]\}$, where $[1] = \{1, 3, 5, \dots\}$ and $[2] = \{2, 4, 6, \dots\}$. Since $2 \in V(p)$, $[2] \in V^*(p)$; since $1 \notin V(p)$, $[1] \notin V^*(p)$. So $V^*(p) = \{[2]\}$.

Any filtration based on W^* must have an accessibility relation that includes $\langle [1], [2]\rangle, \langle [2], [1]\rangle$: since $R12$, we must have $R^*[1][2]$ by Definition fil.4(2a), and since $R23$ we must have $R^*[2][3]$, and $[3] = [1]$. It cannot include $\langle [1], [1]\rangle$: if it did, we'd have $R^*[1][1]$, $\mathfrak{M}, 1 \Vdash \Box p$ but $\mathfrak{M}, 1 \nVdash p$, contradicting (2b). Nothing requires or rules out that $R^*[2][2]$. So, there are two possible filtrations of $\mathfrak{M}$, corresponding to the two accessibility relations

$$\{\langle [1], [2]\rangle, \langle [2], [1]\rangle\} \text{ and } \{\langle [1], [2]\rangle, \langle [2], [1]\rangle, \langle [2], [2]\rangle\}.$$

In either case, p and $\Box p \to p$ are false and $\Box p$ is true at $[1]$; p and $\Box p \to p$ are true and $\Box p$ is false at $[2]$.

Problem fil.3. Consider the following model $\mathfrak{M} = \langle W, R, V\rangle$ where $W = \{0\sigma : \sigma \in \mathbb{B}^*\}$, the set of sequences of 0s and 1s starting with 0, with $R\sigma\sigma'$ iff $\sigma' = \sigma 0$ or $\sigma' = \sigma 1$, and $V(p) = \{\sigma 0 : \sigma \in \mathbb{B}^*\}$ and $V(q) = \{\sigma 1 : \sigma \in \mathbb{B}^* \setminus \{1\}\}$. Here's a picture:

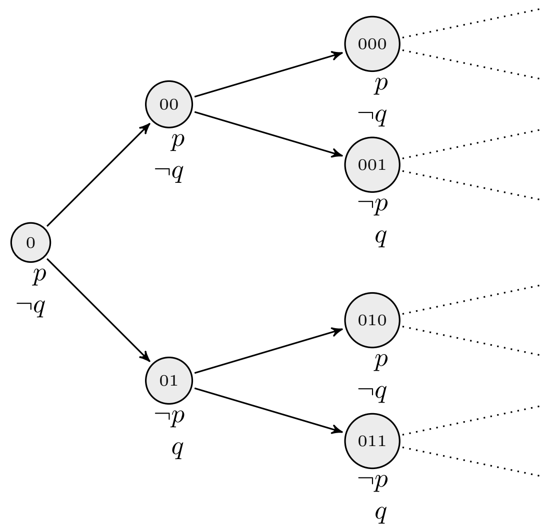

We have $\mathfrak{M}, w \nVdash \Box(p \lor q) \to (\Box p \lor \Box q)$ for every w.

Let Γ be the set of sub-formulas of $\Box(p \lor q) \to (\Box p \lor \Box q)$. What are W^* and V^*? What is the accessibility relation of the finest filtration of $\mathfrak{M}$? Of the coarsest?

fil.5 Filtrations are Finite

We've defined filtrations for any set Γ that is closed under sub-formulas. Nothing in the definition itself guarantees that filtrations are finite. In fact, when Γ is infinite (e.g., is the set of all formulas), it may well be infinite. However, if Γ is finite (e.g., when it is the set of sub-formulas of a given formula φ), so is any filtration through Γ. nml:fil:fin:sec

Proposition fil.12. *If Γ is finite then any filtration $\mathfrak{M}^*$ of a model $\mathfrak{M}$ through Γ is also finite.* *nml:fil:fin:prop:filt-are-finite*

Proof. The size of W^* is the number of different classes $[w]$ under the equivalence relation $\equiv$. Any two worlds u, v in such class—that is, any u and v such that $u \equiv v$—agree on all formulas φ in Γ, $\varphi \in \Gamma$ either φ is true at both u and v, or at neither. So each class $[w]$ corresponds to subset of Γ, namely the set of all $\varphi \in \Gamma$ such that φ is true at the worlds in $[w]$. No two different classes $[u]$ and $[v]$ correspond to the same subset of Γ. For if the set of formulas true at u and that of formulas true at v are the same, then u and v agree on all formulas in Γ, i.e., $u \equiv v$. But then $[u] = [v]$. So, there is an injective function from W^* to $\wp(\Gamma)$, and hence $|W^*| \leq |\wp(\Gamma)|$. Hence if Γ contains n sentences, the cardinality of W^* is no greater than 2^n. □

fil.6 K and S5 have the Finite Model Property

nml:fil:fmp: sec

Definition fil.13. A system Σ of modal logic is said to have the *finite model property* if whenever a formula φ is true at a world in a model of Σ then φ is true at a world in a *finite* model of Σ.

nml:fil:fmp: prop:K-fmp

Proposition fil.14. **K** *has the finite model property.*

Proof. **K** is the set of valid formulas, i.e., any model is a model of **K**. By Theorem fil.5, if $\mathfrak{M}, w \Vdash \varphi$, then $\mathfrak{M}^*, w \Vdash \varphi$ for any filtration of $\mathfrak{M}$ through the set Γ of sub-formulas of φ. Any formula only has finitely many sub-formulas, so Γ is finite. By Proposition fil.12, $|W^*| \leq 2^n$, where n is the number of formulas in Γ. And since **K** imposes no restriction on models, $\mathfrak{M}^*$ is a **K**-model. □

To show that a logic **L** has the finite model property via filtrations it is essential that the filtration of an **L**-model is itself a **L**-model. Often this requires a fair bit of work, and not any filtration yields a **L**-model. However, for universal models, this still holds.

nml:fil:fmp: prop:univ-fin

Proposition fil.15. *Let $\mathcal{U}$ be the class of universal models (see* **??***) and $\mathcal{U}_{\mathrm{Fin}}$ the class of all finite universal models. Then any formula φ is valid in $\mathcal{U}$ if and only if it is valid in $\mathcal{U}_{\mathrm{Fin}}$.*

Proof. Finite universal models are universal models, so the left-to-right direction is trivial. For the right-to left direction, suppose that φ is false at some world w in a universal model $\mathfrak{M}$. Let Γ contain φ as well as all of its subformulas; clearly Γ is finite. Take a filtration $\mathfrak{M}^*$ of $\mathfrak{M}$; then $\mathfrak{M}^*$ is finite by Proposition fil.12, and by Theorem fil.5, φ is false at $[w]$ in $\mathfrak{M}^*$. It remains to observe that $\mathfrak{M}^*$ is also universal: given u and v, by hypothesis Ruv and by Definition fil.4(2), also $R^*[u][v]$. □

nml:fil:fmp: cor:S5fmp

Corollary fil.16. **S5** *has the finite model property.*

Proof. By **??**, if φ is true at a world in some reflexive and euclidean model then it is true at a world in a universal model. By Proposition fil.15, it is true at a world in a finite universal model (namely the filtration of the model through the set of sub-formulas of φ). Every universal model is also reflexive and euclidean; so φ is true at a world in a finite reflexive euclidean model. □

Problem fil.4. Show that any filtration of a serial or reflexive model is also serial or reflexive (respectively).

Problem fil.5. Find a non-symmetric (non-transitive, non-euclidean) filtration of a symmetric (transitive, euclidean) model.

fil.7 S5 is Decidable

The finite model property gives us an easy way to show that systems of modal logic given by schemas are *decidable* (i.e., that there is a computable procedure to determine whether a formulas is derivable in the system or not). nml:fil:dec:sec

Theorem fil.17. **S5** *is decidable.*

Proof. Let φ be given, and suppose the propositional variables occurring in φ are among $p_1, \ldots, p_k$. Since for each n there are only finitely many models with n worlds assigning a value to $p_1, \ldots, p_k$, we can enumerate, *in parallel*, all the theorems of **S5** by generating proofs in some systematic way; and all the models containing 1, 2, ... worlds and checking whether φ fails at a world in some such model. Eventually one of the two parallel processes will give an answer, as by **??** and Corollary fil.16, either φ is derivable or it fails in a finite universal model. □

The above proof works for **S5** because filtrations of universal models are automatically universal. The same holds for reflexivity and seriality, but more work is needed for other properties.

Problem fil.6. Show that any filtration of a serial or reflexive model is also serial or reflexive (respectively).

Problem fil.7. Find a non-symmetric (non-transitive, non-euclidean) filtration of a symmetric (transitive, euclidean) model.

fil.8 Filtrations and Properties of Accessibility

As noted, filtrations of universal, serial, and reflexive models are always also universal, serial, or reflexive. But not every filtration of a symmetric or transitive model is symmetric or transitive, respectively. In some cases, however, it is possible to define filtrations so that this does hold. In order to do so, we proceed as in the definition of the coarsest filtration, but add additional conditions to the definition of R^*. Let Γ be closed under sub-formulas. Consider the relations $C_i(u, v)$ in table fil.1 between worlds u, v in a model $\mathfrak{M} = \langle W, R, V\rangle$. We can define $R^*[u][v]$ on the basis of combinations of these conditions. For instance, if we stipulate that $R^*[u][v]$ iff the condition $C_1(u, v)$ holds, we get exactly the coarsest filtration. If we stipulate $R^*[u][v]$ iff both $C_1(u, v)$ and $C_2(u, v)$ hold, we get a different filtration. It is "finer" than the coarsest since fewer pairs of worlds satisfy $C_1(u, v)$ and $C_2(u, v)$ than $C_1(u, v)$ alone. nml:fil:acc:sec

Theorem fil.18. *Let $\mathfrak{M} = \langle W, R, P\rangle$ be a model, Γ closed under sub-formulas. Let W^* and V^* be defined as in Definition fil.4. Then:* nml:fil:acc:thm:more-filtrations

1. *Suppose $R^*[u][v]$ if and only if $C_1(u, v) \wedge C_2(u, v)$. Then R^* is symmetric, and $\mathfrak{M}^* = \langle W^*, R^*, V^*\rangle$ is a filtration if $\mathfrak{M}$ is symmetric.*

$C_1(u,v)$:	if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Box\varphi$ then $\mathfrak{M}, v \Vdash \varphi$; and if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \varphi$ then $\mathfrak{M}, u \Vdash \Diamond\varphi$;
$C_2(u,v)$:	if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \Box\varphi$ then $\mathfrak{M}, u \Vdash \varphi$; and if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \varphi$ then $\mathfrak{M}, v \Vdash \Diamond\varphi$;
$C_3(u,v)$:	if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Box\varphi$ then $\mathfrak{M}, v \Vdash \Box\varphi$; and if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \Diamond\varphi$ then $\mathfrak{M}, u \Vdash \Diamond\varphi$;
$C_4(u,v)$:	if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \Box\varphi$ then $\mathfrak{M}, u \Vdash \Box\varphi$; and if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Diamond\varphi$ then $\mathfrak{M}, v \Vdash \Diamond\varphi$;

Table fil.1: Conditions on possible worlds for defining filtrations.

nml:fil:acc:
tab:Cn-filtrations

2. *Suppose $R^*[u][v]$ if and only if $C_1(u,v) \wedge C_3(u,v)$. Then R^* is transitive, and $\mathfrak{M}^* = \langle W^*, R^*, V^* \rangle$ is a filtration if $\mathfrak{M}$ is transitive.*

3. *Suppose $R^*[u][v]$ if and only if $C_1(u,v) \wedge C_2(u,v) \wedge C_3(u,v) \wedge C_4(u,v)$. Then R^* is symmetric and transitive, and $\mathfrak{M}^* = \langle W^*, R^*, V^* \rangle$ is a filtration if $\mathfrak{M}$ is symmetric and transitive.*

4. *Suppose R^* is defined as $R^*[u][v]$ if and only if $C_1(u,v) \wedge C_3(u,v) \wedge C_4(u,v)$. Then R^* is transitive and euclidean, and $\mathfrak{M}^* = \langle W^*, R^*, V^* \rangle$ is a filtration if $\mathfrak{M}$ is transitive and euclidean.*

Proof. 1. It's immediate that R^* is symmetric, since $C_1(u,v) \Leftrightarrow C_2(v,u)$ and $C_2(u,v) \Leftrightarrow C_1(v,u)$. So it's left to show that if $\mathfrak{M}$ is symmetric then $\mathfrak{M}^*$ is a filtration through Γ. Condition $C_1(u,v)$ guarantees that (2b) and (2c) of Definition fil.4 are satisfied. So we just have to verify Definition fil.4(2a), i.e., that Ruv implies $R^*[u][v]$.

So suppose Ruv. To show $R^*[u][v]$ we need to establish that $C_1(u,v)$ and $C_2(u,v)$. For C_1: if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \Box\varphi$ then also $\mathfrak{M}, v \Vdash \varphi$ (since Ruv). Similarly, if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \varphi$ then $\mathfrak{M}, u \Vdash \Diamond\varphi$ since Ruv. For C_2: if $\Box\varphi \in \Gamma$ and $\mathfrak{M}, v \Vdash \Box\varphi$ then Ruv implies Rvu by symmetry, so that $\mathfrak{M}, u \Vdash \varphi$. Similarly, if $\Diamond\varphi \in \Gamma$ and $\mathfrak{M}, u \Vdash \varphi$ then $\mathfrak{M}, v \Vdash \Diamond\varphi$ (since Rvu by symmetry).

2. Exercise.

3. Exercise.

4. Exercise. □

Problem fil.8. Complete the proof of Theorem fil.18.

fil.9 Filtrations of Euclidean Models

nml:fil:euc:
sec

The approach of section fil.8 does not work in the case of models that are euclidean or serial and euclidean. Consider the model at the top of Figure fil.2, which is both euclidean and serial. Let $\Gamma = \{p, \Box p\}$. When taking a filtration through Γ, then $[w_1] = [w_3]$ since w_1 and w_3 are the only worlds that agree on Γ. Any filtration will also have the arrow inherited from $\mathfrak{M}$, as depicted in Figure fil.3. That model isn't euclidean. Moreover, we cannot add arrows to that model in order to make it euclidean. We would have to add double arrows between $[w_2]$ and $[w_4]$, and then also between w_2 and w_5. But $\Box p$ is supposed to be true at w_2, while p is false at w_5.

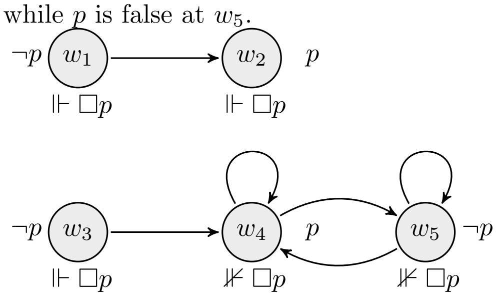

Figure fil.2: A serial and euclidean model. nml:fil:euc: fig:ser-eucl

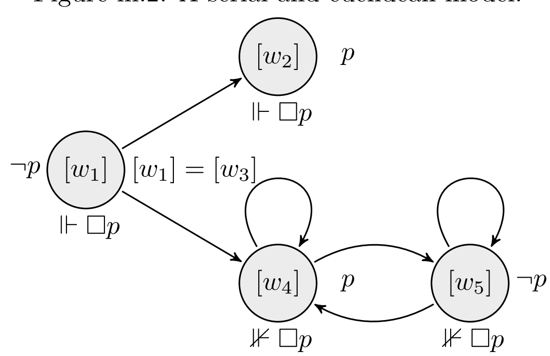

Figure fil.3: The filtration of the model in Figure fil.2.

In particular, to obtain a euclidean flitration it is not enough to consider filtrations through arbitrary Γ's closed under sub-formulas. Instead we need to consider sets Γ that are *modally closed* (see Definition fil.1). Such sets of sentences are infinite, and therefore do not immediately yield a finite model property or the decidability of the corresponding system. nml:fil:euc: fig:ser-eucl2

Theorem fil.19. *Let Γ be modally closed, $\mathfrak{M} = \langle W, R, V\rangle$, and $\mathfrak{M}^* = \langle W^*, R^*, V^*\rangle$ be a coarsest filtration of $\mathfrak{M}$.* nml:fil:euc: thm:modal-closed-filt

1. *If $\mathfrak{M}$ is symmetric, so is $\mathfrak{M}^*$.*
2. *If $\mathfrak{M}$ is transitive, so is $\mathfrak{M}^*$.*
3. *If $\mathfrak{M}$ is euclidean, so is $\mathfrak{M}^*$.*

Proof. 1. If $\mathfrak{M}^*$ is a coarsest filtration, then by definition $R^*[u][v]$ holds if and only if $C_1(u, v)$. For transitivity, suppose $C_1(u, v)$ and $C_1(v, w)$; we have to show $C_1(u, w)$. Suppose $\mathfrak{M}, u \Vdash \Box\varphi$; then $\mathfrak{M}, u \Vdash \Box\Box\varphi$ since 4 is valid in all transitive models; since $\Box\Box\varphi \in \Gamma$ by closure, also by $C_1(u, v)$, $\mathfrak{M}, v \Vdash \Box\varphi$ and by $C_1(v, w)$, also $\mathfrak{M}, w \Vdash \varphi$. Suppose $\mathfrak{M}, w \Vdash \varphi$; then $\mathfrak{M}, v \Vdash \Diamond\varphi$ by $C_1(v, w)$, since $\Diamond\varphi \in \Gamma$ by modal closure. By $C_1(u, v)$, we get $\mathfrak{M}, u \Vdash \Diamond\Diamond\varphi$ since $\Diamond\Diamond\varphi \in \Gamma$ by modal closure. Since $4_\Diamond$ is valid in all transitive models, $\mathfrak{M}, u \Vdash \Diamond\varphi$.

2. Exercise. Use the fact that both 5 and $5_\Diamond$ are valid in all euclidean models.

3. Exercise. Use the fact that B and $B_\Diamond$ are valid in all symmetric models. □

Problem fil.9. Complete the proof of Theorem fil.19.

Photo Credits

Bibliography