

Requirements, specifications, and minimal refinement

Nikos Gorogiannis Mark Ryan

2002

Abstract

Refinement is usually employed to produce more concrete versions of a specification, or to add new requirements to it. However, during specification revision one may over-refine, thus incorporating unnecessary requirements. In this paper, we argue that this process can be formalised by the notion of minimal refinements, hence avoiding over-refinement, and prove that this definition is well-behaved theoretically as well as computationally.

1 Introduction

A common way to develop computer systems (whether hardware or software) is by *refinement*: one starts with an abstract specification, and refines it gradually (e.g. [1]). The refinements may be triggered by the need to satisfy additional requirements. However, one would like to avoid refining the specification too much, in order to keep it flexible and avoid building in unnecessary assumptions. In this paper, we study *minimal refinements*. We address the question: given a specification and a requirement, what is the smallest refinement of the specification which will make it satisfy the requirement?

Example 1. *A university department has a policy which governs access to student marks, including perhaps the requirements:*

- *a student has read-access to all his marks;*
- *a student does not have write-access to any of his marks;*
- *a professor has read-access to all student marks, and write-access to the marks of the modules she teaches.*

This system allows certain access operations and denies others. Because it is under-specified it may also be non-deterministic about the outcome (whether access allowed or not) of some operations. The policy may be encoded as a non-deterministic transition system M . Any implementation which satisfies the requirements, i.e. which refines M , is acceptable.

A further requirement is now imposed upon the department, perhaps by new legislation, such as:

- *no student may have read-access to another student's marks.*

To incorporate this, we seek a refinement of M which satisfies the new requirement. Naturally, we do not want to refine too much, unnecessarily losing flexibility with respect of further requirements, so we attempt to refine minimally, just enough to satisfy the new constraint.

We assume that specifications and implementations of systems are represented as models (Kripke models or transition systems), and that requirements are presented as logical formulas. We study the operator $*$ which takes a model M and a formula ϕ , and returns a set of models $M * \phi$ which is the set of least refinements of M which satisfy ϕ . We explore the following properties of this operation:

- When do minimal refinements exist?
- When are the properties of the minimal refinements decidable?

These questions are studied in a variety of contexts, such as: finite models, serial models, and m-saturated models.

We will use Kripke models to model systems. A well-understood notion of refinement in this context is simulation [12]. It has been studied extensively as well as having served as the basis for a multitude of other more fine-grained proposed notions for refinement. The logic used in the following is the polymodal logic K_n (essentially Hennessy-Milner logic [8] extended with propositional information) with an outlook on temporal logics. Using simulation, we define an ordering that depends on M , which captures refinement of M and that is related to the simulation preorder (see e.g. [3, 7]). Then, $M * \phi$ is defined as the set of the models of ϕ that are minimal with respect to the ordering.

Below we prove that for two important classes of Kripke models, the m-saturated and the finite models, the operation is well-behaved. We characterise the conditions that the specification and the property need to satisfy in order for the operation to yield non-trivial results (i.e. not just all the models of ϕ). In the case of finite models, we prove that checking whether an implementation is minimal is decidable and that for properties in a fragment of K_n (and also for a fragment of CTL*), checking whether such properties are true on the results of the operation is decidable.

The definition of $M * \phi$ is reminiscent of, and indeed inspired from, theory change and non-monotonic reasoning. In those fields one of the ways to define a theory change operation is to define an ordering on possible worlds that captures a notion of closeness to the initial world, and then minimising with respect to that ordering within a prescribed set of worlds. In this sense $M * \phi$ is a non-monotonic operation since it may be the case that $M * \phi \not\models \psi$ while $M \models \psi$.

2 Definitions

Let \mathcal{A} be a set of l atomic propositions. The modal language \mathcal{L} of the logic K_n on \mathcal{A} with k modalities is defined inductively

- if $p \in \mathcal{A}$ then $p \in \mathcal{L}$,
- if $\phi, \psi \in \mathcal{L}$ then $\neg\phi, \phi \wedge \psi \in \mathcal{L}$,
- if $\phi \in \mathcal{L}$ then $\diamond_i\phi \in \mathcal{L}$ for all $1 \leq i \leq k$.

The usual propositional abbreviations apply as well as the modal $\Box_i \equiv \neg\diamond_i\neg$. The degree $\text{deg}(\phi)$ of a formula ϕ is defined as the maximum nesting depth of modalities in ϕ .

A *Kripke model* M for \mathcal{L} is a tuple $\langle W_M, r_M, R_M^1, \dots, R_M^k, v_M \rangle$. W_M is a set of *states* or *worlds*. r_M is a distinguished state in W_M called the *initial state* or the *root*. $R_M^i \subseteq W_M \times W_M$ are *accessibility relations* and $v_M : W_M \rightarrow 2^{\mathcal{A}}$ is a *valuation* for the propositional letters. Satisfaction of formulas at a state s is defined inductively by the usual propositional clauses along with the modal one: $M, s \models \diamond_i\phi$ iff there exists a state $t \in W_M$ such that $(s, t) \in R_M^i$ and $M, t \models \phi$. We will write $s \models \phi$ when the model is obvious. By $|M|$ we denote the cardinality of W_M . A model M is finite iff $|M|$ is finite.

A *path* is a finite sequence of states such that for any pair of states s_i, s_{i+1} in the sequence, there exists a j such that $(s_i, s_{i+1}) \in R_M^j$. The *depth* of a state s is defined as the minimum length of a path from the root to s if such a path exists, otherwise as ω .

A model M for a logic with a single modality is called *serial* if the single accessibility relation R_M is serial, i.e. iff for all states $s \in W_M$ there exists a state $t \in W_M$ such that $(s, t) \in R_M$.

The set of sentences true at a state s is denoted by $\text{th}(s)$. In the following we will focus on validity of formulas on the root and not on the whole model as is usual in modal logic. This approach is commonplace in the temporal logic literature where models represent transition systems with a starting state. Thus we define the theory of a model to be the theory of its root, $\text{th}(M) = \text{th}(r_M)$. Since the root is our ‘entry point’ in a model, we will only consider models whose states are all reachable from the root. Two models M, N are *logically equivalent* iff $\text{th}(M) = \text{th}(N)$.

The logic usually studied in modal logic is the one enforced by global validity on frames. In other words, $\Gamma \models \phi$ is taken to mean that for all frames \mathcal{F} , if $\mathcal{F} \models \Gamma$, then $\mathcal{F} \models \phi$. As noted above, we employ a local entailment relation at the level of states of models, i.e. taking $\Gamma \models \phi$ to mean that for all models M and all states $s \in W_M$, if $M, s \models \Gamma$ then $M, s \models \phi$. These two definitions give rise to the same logic, a fact witnessed by the strong completeness of K_n (see e.g. [2]). To simplify our exposition, we will use an axiomatisation that is equivalent to the usual for K_n but validates the deduction theorem at the cost of losing the necessitation rule. This axiomatisation has modus ponens as its sole rule of inference and as axioms it has all propositional tautologies, possibly

prefixed by an arbitrary sequence of box modalities and any formula of the form $\Box_{i_1} \dots \Box_{i_n} (\Box_j (\phi \Rightarrow \psi) \Rightarrow (\Box_j \phi \Rightarrow \Box_j \psi))$.

Let M, N be models and $B \subseteq W_M \times W_N$ a relation. B is a *bisimulation* if

- It relates the initial states, $(r_M, r_N) \in B$,
- It respects the valuations, $(s, t) \in B$ implies $v_M(s) = v_N(t)$,
- If $(s, t) \in B$ and s' is an R_M^j -successor of s then there exists t' , an R_N^j -successor of t , such that $(s', t') \in B$, for all j (the *forth* condition),
- If $(s, t) \in B$ and t' is an R_N^j -successor of t then there exists s' , an R_M^j -successor of s , such that $(s', t') \in B$, for all j (the *back* condition).

If there exists a bisimulation between M, N then M and N are *bisimilar*, written $M \sim N$ and it follows that $\text{th}(M) = \text{th}(N)$.

An approximation of bisimulation is *n-bisimulation*. Two models M, N are *n-bisimilar*, written $M \sim_n N$ iff there exists a sequence of relations $\sim_n \subseteq \dots \subseteq \sim_0 \subseteq W_M \times W_N$ such that

- $r_M \sim_n r_N$,
- For all $1 \leq i \leq k$ and all $m < n$, if $s \sim_{m+1} t$ and s' is an R_M^i -successor of s then there is an R_N^i -successor t' of t such that $s' \sim_m t'$,
- For all $1 \leq i \leq k$ and all $m < n$, if $s \sim_{m+1} t$ and t' is an R_N^i -successor of t then there is an R_M^i -successor s' of s such that $s' \sim_m t'$,
- For all $m \leq n$, if $s \sim_m t$ then $v_M(s) = v_N(t)$.

Bisimilarity implies *n-bisimilarity* for all n , but the converse is not true in general. Another standard result about *n-bisimulations* is that $M \sim_n N$ iff for all formulas ϕ with $\text{deg}(\phi) \leq n$, $M \models \phi$ iff $N \models \phi$. Also, a result which we will make use of below is that for all n there is an effective procedure for computing a finite set of finite models \mathcal{T}_n such that (a) every model in \mathcal{T}_n is a tree of depth at most n and (b) for any model M there is a tree $T \in \mathcal{T}_n$ such that $M \sim_n T$. These results can be found in [13].

A formula is called *positive universal* iff it is made up only from $p, \neg p, \wedge, \vee$ and \Box_i for all $1 \leq i \leq k$. \mathcal{L}_{PU} is the subset of \mathcal{L} that consists of positive universal formulas. If s is a state then $\text{PU}(s) = \mathcal{L}_{\text{PU}} \cap \text{th}(s)$. If M is a model, then $\text{PU}(M) = \text{PU}(r_M)$. Dually, a *positive existential* formula is made up from $p, \neg p, \wedge, \vee$ and \Diamond_i . \mathcal{L}_{PE} and PE are defined similarly and are duals of \mathcal{L}_{PU} and PU respectively. Note that the negation of a PU formula is a PE one and vice versa. If P is a set of PU sentences then P^c is the complement of P with respect to \mathcal{L}_{PU} . \overline{P} contains the negation of every formula in P .

Intuitively, positive universal formulas describe restrictions on what states are accessible. In the context of transition systems, PU formulas prescribe what conditions a sequence of actions must satisfy if it is to be allowed. Dually, a PE formula asserts the possibility of the execution of a sequence of actions.

Let M be a model and $s \in W_M$ a state. A set of sentences T will be called *satisfiable on the successors of s* iff for each relation R_M^i there exists a state $t \in W_M$ such that $(s, t) \in R_M^i$ and $T \subseteq \text{th}(t)$. Similarly, T will be called *finitely-satisfiable on the successors of s* iff for each relation R_M^i and for any finite set of sentences $F \subseteq T$ there exists an R_M^i -successor t of s such that $F \subseteq \text{th}(t)$. A state s is called *m-saturated* iff for any set of sentences T , if T is finitely-satisfiable on the successors of s , then it is satisfiable on the successors of s . A model is m-saturated if all its states are m-saturated. $\text{mod}_m(\phi)$ is the class of m-saturated models M of ϕ . We write MSAT for the class of m-saturated models. Notice that MSAT is bisimulation-closed.

In the following we will use the *ultrafilter extension* of a model. We will not make reference to the internals of the construction, just to two of its properties: the ultrafilter extension of a model M is another model $\text{ue}(M)$ that is logically equivalent to M and also, $\text{ue}(M)$ is m-saturated. Accounts of the construction appear in many places, e.g. [2].

A class of models has the *Hennessey-Milner property* whenever for every pair of its models, they are bisimilar iff they are logically equivalent. In other words, models in a Hennessey-Milner class are completely characterised by the logic, i.e. if two such models are not bisimilar then there is a witnessing formula that distinguishes them.

MSAT has the following important properties [9]

- It subsumes the class of image-finite models (and hence the finite ones).
- It has the Hennessey-Milner property.
- It is maximal in the sense that no proper superclass of MSAT has the Hennessey-Milner property.
- It has also been used to provide semantics for process algebras.

Let M, N be models and $S \subseteq W_M \times W_N$ a relation on their states. S will be called a *simulation* iff it satisfies the first three clauses in the definition of bisimulation, i.e. it must link the initial states, preserve valuations and respect the accessibility relations but in one-way only (the forth condition). If there exists a simulation from M to N we write $M \rightarrow N$ or $N \leftarrow M$ and say that N simulates M or that M is simulated by N . Whenever $M \leftarrow N$ and $M \rightarrow N$ we will say that M and N are *similar* or *simulation equivalent* and write $M \rightleftharpoons N$. It is easy to check that simulations are transitive.

Let \mathcal{M} be a class of models. An ordering \leq over \mathcal{M} is *stoppered for a formula ϕ* iff for any model $M \in \text{mod}_{\mathcal{M}}(\phi)$ there is another model $N \in \text{mod}_{\mathcal{M}}(\phi)$ such that $N \leq M$ and that N is \leq -minimal in $\text{mod}_{\mathcal{M}}(\phi)$. The definition is extended for sets of sentences in the obvious way.

3 Results

Let M, N_1, N_2 be models. We define an ordering \leq_M such that $N_1 \leq_M N_2$ iff

1. $M \leftarrow N_1 \leftarrow N_2$ or
2. $M \leftarrow N_1$ but $M \not\leftarrow N_2$ or
3. $N_1 \rightleftharpoons N_2$.

It is not hard to prove that this ordering is transitive and reflexive. By taking similarity as the main equivalence notion between models, antisymmetry is obtained, i.e. if $A \leq_M B$ and $B \leq_M A$ then $A \rightleftharpoons B$. In other words, \leq_M is a partial order.

Let \mathcal{M} be a class of models, M a model in \mathcal{M} and T a set of sentences. We define an operation $*_{\mathcal{M}} : \mathcal{M} \times 2^{\mathcal{L}} \rightarrow 2^{\mathcal{M}}$

$$M *_{\mathcal{M}} T = \min_{\leq_M}(\text{mod}_{\mathcal{M}}(T))$$

This definition reminds one of a type of theory change which is known as *update* [11]. It is a point-wise definition, i.e. the ordering depends on a model rather than an arbitrary theory as is usual in the case of *revisions*, the other well-known type of theory change (see, e.g. [6, 10]). In addition, the ordering is partial, a condition which automatically validates the update axioms via the representation theorem mentioned in [11].

Given such an operation, several questions arise. Firstly, it is not obvious that it is well-defined, i.e. whether the existence of minimal models is guaranteed so that $M *_{\mathcal{M}} T \neq \emptyset$. We address this question in propositions 1 and 2, for the class of m-saturated models and arbitrary sets of sentences and for the class of finite models and arbitrary sentences, respectively.

Moreover, it is of interest to know the conditions that guarantee non-triviality of the results of the operation, or in other words, when it is the case that $M *_{\mathcal{M}} T \subset \text{mod}_{\mathcal{M}}(T)$. The necessary and sufficient conditions for non-triviality are presented in lemmas 4 and 6 for m-saturated and finite models, respectively. In addition, the decidability of determining non-triviality for a finite model M and a formula ϕ is proved in lemma 7.

Finally, in the case of finite models, we prove that two interesting problems are decidable: firstly, that checking minimality of a finite model N with respect to a finite model M and a formula ϕ is decidable (lemma 10). Secondly, that reasoning within a fragment of the language about the results of the operation is decidable, i.e. answering queries of the form $M * \phi \models \psi$ (proposition 4).

The first three lemmas characterise simulation in syntactic terms, and establish an exact match in the m-saturated case.

Lemma 1. (*Folklore*) *If M, N are models such that $M \leftarrow N$, then $\text{PU}(M) \subseteq \text{PU}(N)$.*

Lemma 2. (*Folklore*) *Let M, N be models. If $\text{PU}(M) \subseteq \text{PU}(N)$ and M is m-saturated, then there exists a simulation from N to M , $M \leftarrow N$.*

Proof. For convenience we will work with PE formulas, the dual of PU ones. Note that $\text{PU}(s) \subseteq \text{PU}(t)$ iff $\text{PE}(s) \supseteq \text{PE}(t)$. Define a relation S such that

$(s, t) \in S$ iff $s \in W_N$, $t \in W_M$ and $\text{PE}(s) \subseteq \text{PE}(t)$. We prove that S is a simulation. Obviously it respects the valuations, i.e. if $(s, t) \in S$ then $v_N(s) = v_M(t)$. Assume that s has a successor s' with respect to a relation R_N^i . Let P be the set of PE sentences of s' . For any finite subset $F \subset P$, $s' \models \bigwedge F$ and thus $s \models \diamond_i \bigwedge F$. $\diamond_i \bigwedge F$ is a PE formula, so by definition it is satisfied at t . Thus there is an R_M^i -successor of t that satisfies $\bigwedge F$. In other words, P is finitely-satisfiable on the successors of t . M however is m-saturated, thus there is an R_M^i -successor t' of t that satisfies P and as such $\text{PE}(s') \subseteq \text{PE}(t')$.

So, S is a simulation whenever it is non-empty and it relates the initial states. Those conditions are satisfied by the assumption $\text{PU}(M) \subseteq \text{PU}(N)$ or equivalently $\text{PE}(N) \subseteq \text{PE}(M)$. \square

Let T be a set of sentences. T is *closed under taking disjuncts* iff whenever $\phi \vee \psi \in T$ then $\phi \in T$ or $\psi \in T$. T is *closed under \mathcal{L}_{PU} -consequence* iff whenever $T \vdash \phi$ and $\phi \in \mathcal{L}_{\text{PU}}$ then $\phi \in T$.

Lemma 3. *Let $P \subseteq \mathcal{L}_{\text{PU}}$. There exists a model M such that $P = \text{PU}(M)$ iff P is consistent, closed under \mathcal{L}_{PU} -consequence and taking disjuncts.*

Proof. The left-to-right direction is trivial. Right-to-left: for a model M to have exactly P as its set of PU formulas, it must satisfy P and falsify its complement with respect to \mathcal{L}_{PU} . In other words, there exists such a model iff $P, P^c \not\vdash \perp$. Assume the latter is not the case. Then there exist formulas $\phi, \psi_1, \dots, \psi_m$ such that $\phi \in P$ (note that P is closed under conjunction), $\neg\psi_i \in P^c$ and $\phi, \neg\psi_1, \dots, \neg\psi_m \vdash \perp$. But then, $\phi \vdash \psi_1 \vee \dots \vee \psi_m$ and since P is closed under \mathcal{L}_{PU} -consequence, $\psi_1 \vee \dots \vee \psi_m \in P$. P is also closed under taking disjuncts so there exists $1 \leq j \leq m$ such that $\psi_j \in P$ which is a contradiction because $\psi_j \in P^c$. \square

If no model of a set of sentences T is simulated by a model M , then as noted in the beginning of this section, all models of T will be incomparable with respect to the ordering \leq_M , and thus, $M *_m T = \text{mod}_m(T)$. If there is at least one such model in $\text{mod}_m(T)$, then $*_m$ will return a strict subset of $\text{mod}_m(T)$, in view of the second clause of the definition of the ordering. The conditions under which this happens are characterised in the next lemma.

Lemma 4. *Let M be an m-saturated model and T a set of sentences. Then, there exists an m-saturated model N of T such that $M \leftarrow N$ iff $\text{PU}(M), T \not\vdash \perp$.*

Proof. Left-to-right: Since $M \leftarrow N$ it follows from lemma 1 that $\text{PU}(M) \subseteq \text{PU}(N)$. Thus N is a model of both T and $\text{PU}(M)$.

Right-to-left: Let N be a model of $\text{PU}(M) \cup T$. Then, $\text{PU}(M) \subseteq \text{PU}(N)$. Since N may not be m-saturated, we take the ultrafilter-extension of N , $\text{ue}(N)$ which is logically equivalent to N and as such a model of $\text{PU}(M) \cup T$, and m-saturated. It follows that $T \subseteq \text{th}(\text{ue}(N))$ and that $\text{PU}(M) \subseteq \text{PU}(\text{ue}(N))$. As M is m-saturated it follows from lemma 2 that $M \leftarrow \text{ue}(N)$. \square

The following lemma and proposition concern stopperedness of the ordering for m -saturated models. Lemma 5 enables us to apply Zorn's lemma by proving that for any suitable chain (i.e. a totally ordered set of models), a suitable lower bound can be found, and indeed, the infimum.

Lemma 5. *Let M be an m -saturated model and T a consistent set of sentences of which M is not a model. Let $C \subseteq \text{mod}_m(T)$ be a nonempty chain with respect to \leq_M where all of its members are simulated by M . Then there exists an m -saturated model of T which is the infimum of C (modulo simulation equivalence).*

Proof. Define $P = \bigcap_{N \in C} \text{PU}(N)$. Since any model N in the chain is simulated by M , $\text{PU}(M) \subseteq \text{PU}(N)$ and therefore $\text{PU}(M) \subseteq P$. Also, for any two models $A, B \in C$ it will be the case that $\text{PU}(A) \subseteq \text{PU}(B)$ or $\text{PU}(B) \subseteq \text{PU}(A)$. We will prove that there exists a model I with $\text{PU}(I) = P$ which satisfies T . P is obviously consistent as a subset of consistent sets. Also, it is easy to check that P is closed under \mathcal{L}_{PU} -consequence.

We now prove that P is closed under taking disjuncts. Assume $\phi \vee \psi \in P$. Then, for all $L \in C$, $L \models \phi \vee \psi$. If all the models in C satisfy ϕ we are done, so assume that there exists a pair of models $N, N' \in C$ such that $N \models \phi \wedge \neg\psi$ and $N' \models \neg\phi \wedge \psi$. But this contradicts the fact mentioned above, that $\text{PU}(N) \subseteq \text{PU}(N')$ or $\text{PU}(N') \subseteq \text{PU}(N)$. Hence P is closed under taking disjuncts.

From lemma 3 it follows that $P \cup \overline{P^c}$ is consistent. Assume that $P, \overline{P^c}, T \vdash \perp$. Then there exist $\neg\phi_1, \dots, \neg\phi_n \in \overline{P^c}$ such that $P, T, \neg\phi_1, \dots, \neg\phi_n \vdash \perp$ or equivalently $P, T \vdash \phi_1 \vee \dots \vee \phi_n$. Thus, for all $N \in C$, $N \models \phi_1 \vee \dots \vee \phi_n$, hence $\phi_1 \vee \dots \vee \phi_n \in \text{PU}(N)$ and therefore $\phi_1 \vee \dots \vee \phi_n \in P$. As P is closed under taking disjuncts there is one disjunct ϕ_j such that $\phi_j \in P$, which is a contradiction. So there is a model I of $P \cup \overline{P^c} \cup T$. I need not be m -saturated, but its ultrafilter extension $\text{ue}(I)$ is, and as it is logically equivalent to I it will satisfy $P \cup \overline{P^c} \cup T$ too.

By the definition of P we have that for all $N \in C$, $\text{PU}(\text{ue}(I)) \subseteq \text{PU}(N)$. Thus, by lemma 2 we get that $\text{ue}(I) \leftarrow N$. Also, $\text{PU}(M) \subseteq \text{PU}(\text{ue}(I))$ which implies that $M \leftarrow \text{ue}(I)$. So, $\text{ue}(I)$ is a lower bound of C with respect to \leq_M . In addition, for any other lower bound L of C , it follows that $\text{PU}(L) \subseteq \bigcap_{N \in C} \text{PU}(N)$ and thus that $\text{ue}(I)$ is the infimum of C (modulo similarity). \square

In propositions 1 and 2 we prove stopperedness for m -saturated and finite models, respectively. The application of Zorn's lemma is usually a crucial part of such proofs. The commonly cited version of Zorn's lemma, however, is not enough to yield stopperedness when its premises are satisfied. We use an easily derivable, but stronger version: if X is a partially-ordered set and any well-ordered subset of X has a lower bound in X , then for any element of $s \in X$, there exists a minimal element $s' \in X$ that is comparable to s , i.e. $s' \leq s$.

Proposition 1. *Let M be an m -saturated model. The ordering \leq_M over the class of m -saturated models is stoppered for any consistent set of sentences T .*

Proof. If $T \subseteq \text{th}(M)$ then, of course, M is a *minimum* with respect to \leq_M in $\text{mod}_m(T)$, as well as any other m -saturated model N of T such that $M \rightleftharpoons N$. It follows that for any m -saturated model L of T there is an m -saturated model of T , i.e. M , which is minimal and $M \leq_M N$. In the case where $M \notin \text{mod}_m(T)$, it may or may not be the case that $\text{PU}(M) \cup T$ is consistent. If not, then by applying lemma 4 it follows that there are no models in $\text{mod}_m(T)$ that are simulated by M . Hence, only the third clause of the definition of \leq_M can ever apply, rendering all (equivalence classes under simulation of) models in $\text{mod}_m(T)$ incomparable. In this case, for any model $N \in \text{mod}_m(T)$ there is a model N' (namely N itself) such that $N' \leq_M N$, where N' is minimal.

Thus, we assume that $\text{PU}(M) \cup T$ is consistent. Because of the second clause of the definition of the ordering, it is easy to see that in this case the set of minimal elements will be a subset of $\text{mod}_m(\text{PU}(M) \cup T)$. Therefore we restrict our attention to the models in $\text{mod}_m(\text{PU}(M) \cup T)$ which, by virtue of lemma 4, are all simulated by M . Then, for a chain in $\text{mod}_m(\text{PU}(M) \cup T)$, lemma 5 applies. Since it asserts something about any chain, i.e. any totally-ordered set of models, it specialises directly to well-ordered chains of models. Therefore, by Zorn's lemma, for any model $N \in \text{mod}_m(T)$ there exists another model $N' \in \text{mod}_m(T)$ such that N' is minimal and $N' \leq_M N$. \square

This concludes our set of results for m -saturated models. For finite models, we start again from characterising the conditions under which the operation is non-trivial, and also prove the decidability of determining non-triviality. We will use $\text{mod}_f(\phi)$ to denote the class of finite models that satisfy ϕ .

Lemma 6. *Let M be a finite model and ϕ a formula. Then, $\text{PU}(M), \phi \not\vdash \perp$ iff there exists a finite tree L of depth at most $\text{deg}(\phi)$ such that $L \models \phi$ and $M \leftarrow L$.*

Proof. The right-to-left direction is trivial. So, we assume the former and apply lemma 4 to obtain a (possibly infinite) model K such that $K \models \phi$ and $M \leftarrow K$. We construct a finite model L of ϕ such that $K \leftarrow L$. For a fixed n there is a (computable) finite collection of trees \mathcal{T}_n of depth up to n such that for any model A there is a tree $T \in \mathcal{T}_n$ such that $A \sim_n T$. Let $n = \text{deg}(\phi)$. Let L be the tree in \mathcal{T}_n such that $L \sim_n K$. Obviously $L \models \phi$. The n -bisimulation between K and L is also a (backwards) simulation between K and L , i.e. $K \leftarrow L$. Because of transitivity of simulations, $M \leftarrow L$. \square

Lemma 7. *Let M be a finite model and ϕ a formula. The decision problem of whether there exists a finite model L of ϕ such that $M \leftarrow L$ is decidable.*

Proof. From lemma 6 it follows that if there is such a model there is also a finite one. Indeed one with depth at most $n = \text{deg}(\phi)$. We produce \mathcal{T}_n . For each model T in \mathcal{T}_n we check whether $T \models \phi$ and whether $M \leftarrow T$ (both problems are decidable because M and T are finite). \square

Let L be a model, and $s \in W_L$ one of its states. s is said to have *in-degree one* whenever it has a unique ancestor with respect to the union of all accessibility

relations in L . L will be called *smooth* iff every state in W_L apart from the root has in-degree one and finite depth, or in other words, L is a countable tree. For every model L there is a smooth one L^s such that $L \sim L^s$. The proof of this result as well as of a general version of the following lemma can be found in [5]. This lemma will allow us to concentrate on simple simulations, i.e. functional ones, in what follows.

Lemma 8. *Let K and M be models such that K is smooth, M is m -saturated and $M \leftarrow K$. Then there exists a functional simulation from K to M .*

Proof. We define a function $S : W_K \rightarrow W_M$ and prove by induction that for any $t \in W_K$, $\text{PE}(t) \subseteq \text{PE}(S(t))$. We set $S(r_K) = r_M$. Since $M \leftarrow K$ it follows from lemma 1 that $\text{PU}(M) \subseteq \text{PU}(K)$ and thus $\text{PU}(S(r_K)) \subseteq \text{PU}(r_K)$, or $\text{PE}(r_K) \subseteq \text{PE}(S(r_K))$.

Assume that S has been defined for all states in K of depth up to $n - 1$ and let $t \in W_K$ be a state of depth n . Since K is smooth, t has a uniquely defined ancestor t' with respect to some relation R_K^i . By the inductive hypothesis, $\text{PE}(t') \subseteq \text{PE}(S(t'))$. So, for any finite set of PE sentences $F \subseteq \text{PE}(t)$, it follows that $t' \models \diamond_i \wedge F$, hence $S(t') \models \diamond_i \wedge F$, and as such, there exists a $u \in W_M$ such that $u \models \wedge F$ and $(S(t'), u) \in R_M^i$. In other words, $\text{PE}(t)$ is finitely satisfiable on the R_M^i -successors of $S(t')$ which through the m -saturation of M gives us that $\text{PE}(t)$ is satisfiable at a R_M^i -successor u' . We set $S(t) = u'$ and this completes the proof. \square

In the following lemma we construct a model, the set of states of which is defined by the disjoint union of a collection of (sets of states of) models. To that end we use the following notational device: if $\mathcal{W} = \{A, B, \dots\}$ is a family of models then an element of the disjoint union of the sets of states of models in \mathcal{W} is written as $\langle Z, s \rangle$ where Z is a model in \mathcal{W} and s is a state in Z , i.e. $s \in W_Z$. Lemma 9 is the basis for most of the results concerning finite models; it asserts that when $M \leftarrow K$ for some finite model M and a possibly infinite model K , with $K \models \phi$, then there is a finite model of a bounded size that satisfies ϕ and stands in-between M and K .

Lemma 9. *Let M be a finite model and ϕ a sentence. Assume that there exists a (possibly infinite) model K of ϕ such that $M \leftarrow K$. Then there exists a finite model L of ϕ such that $M \leftarrow L \leftarrow K$. In addition, the size of L is bounded by a computable function f dependent on M and ϕ .*

Proof. Let U be the smooth counterpart of K . Since $U \sim K$ and $M \leftarrow K$ it follows that $M \leftarrow U$. Moreover, since M is finite it is also m -saturated thus lemma 8 applies, giving us a functional simulation S between U and M .

Let $n = \text{deg}(\phi)$. Let A be the submodel of U , having the same root and such that no state has depth more than $n - 1$. Formally $W_A = \{s \in W_U \mid \text{depth}(s) \leq n - 1\}$, $r_A = r_U$, $R_A^i = R_U^i \cap W_A \times W_A$ for all $1 \leq i \leq k$ and v_A is the restriction of v_U on W_A .

If $t \in W_U$ then U_t is the generated submodel of U with t as its root. Similarly, by M_s we denote the generated submodel of M with s as its root. It is easy to see that since S is functional, the image of U_t under S is a submodel of $M_{S(t)}$.

Define a model N in the following way:

1. W_N is the disjoint union of W_A , and of $W_{M_{S(t)}}$ for all $t \in W_U$ with $\text{depth}(t) = n$. In symbols, if $t \in W_A$ then $\langle A, t \rangle \in W_N$ and if $t' \in W_{M_{S(t)}}$ for some $t \in W_U$ with $\text{depth}(t) = n$ then $\langle M_{S(t)}, t' \rangle \in W_N$. The latter is well-defined because for any state t' in W_U with depth n or more, from the smoothness of U it follows that there is a unique ancestor of depth n of t' .
2. $r_N = \langle A, r_A \rangle$.
3. R_N^i is the disjoint union of R_A^i and $R_{M_{S(t)}}^i$ for all t of depth n , along with another component: for all states $s \in W_U$ with depth $n - 1$ (and hence in W_A), if for some i , $(s, t) \in R_U^i$ then $(\langle A, s \rangle, \langle M_{S(t)}, t \rangle) \in R_N^i$.
4. v_N is defined in the natural way, i.e. if $\langle A, t \rangle \in W_N$ then $v_N(\langle A, t \rangle) = v_A(t)$. If $\langle M_{S(t)}, t' \rangle \in W_N$ for some suitable t' and t , then $v_N(\langle M_{S(t)}, t' \rangle) = v_{M_{S(t)}}(t')$.

From the definition of N it follows that $N \sim_n K$ and thus, $N \models \phi$.

Define a relation S_{UN} as the smallest one with the following properties

- For all $t \in W_U$ with $\text{depth}(t) < n$ (thus in W_A too), $(t, \langle A, t \rangle) \in S_{UN}$.
- If $(s, s') \in S$ such that there is a state $t \in W_U$ with depth n such that $s \in W_{U_t}$, then $(s, \langle M_{S(t)}, s' \rangle) \in S_{UN}$.

Similarly, define S_{NM}

- If $(s, t) \in S$ where $\text{depth}(s) < n$ then $(\langle A, s \rangle, t) \in S_{NM}$.
- For all $\langle M_{S(t)}, t' \rangle \in W_N$, $(\langle M_{S(t)}, t' \rangle, t') \in S_{NM}$.

It is easy but tedious to verify that S_{UN} and S_{NM} are simulations. Thus, $M \leftarrow N \leftarrow K$.

We now prove that N has a finite bisimilar counterpart L . Since the states at depth n are all initial states of generated submodels of M , there can be at most $|M|$ non-bisimilar ones. The nodes at depth $n - 1$ can have 2^l different propositional valuations where l is the number of atomic propositions. Also, a node at depth $n - 1$ can have $2^{|M|}$ possible different combinations of children from depth n , so the maximum number of non-bisimilar states at depth $n - 1$ is $2^l \cdot 2^{|M| \cdot k}$, where k is the number of accessibility relations. In general, if there are $g(i + 1)$ non-bisimilar states at depth $i + 1$, there are $g(i) = 2^{l + g(i + 1) \cdot k}$ many non-bisimilar states at depth i . Thus, the total number of states will consist of (a) the initial state, (b) the sum of the number of states at each layer, with depth ranging from 1 to $n - 1$, and (c) the number of non-bisimilar states in all

the possible generated submodels of M , i.e. $|M|^2$. So, there is a finite model L with at most $f(M, \phi) = 1 + \sum_{i=1}^{\deg(\phi)-1} g(i) + |M|^2$ states, which is bisimilar to N .

Since $L \sim N$ and $M \leftarrow N \leftarrow K$ it is easy to see that $M \leftarrow L \leftarrow K$ and that $L \models \phi$. \square

Proposition 2. *Let M be a finite model. The ordering \leq_M over the class of finite models is stoppered for any consistent sentence ϕ .*

Proof. As in the proof of proposition 1, it is easy to check that when $M \models \phi$ or $\text{PU}(M), \phi \vdash \perp$ then for any model $N \in \text{mod}_f(\phi)$ there exists a model $N' \in \text{mod}_f(\phi)$ such that $N' \leq_M N$ and N' is minimal. So we assume that $M \not\models \phi$, that $\text{PU}(M), \phi \not\vdash \perp$ and restrict our attention to the models in $\text{mod}_f(\text{PU}(M) \cup \{\phi\})$.

Let $C \subseteq \text{mod}_f(\text{PU}(M) \cup \{\phi\})$ be a chain with respect to \leq_M . Since finite models are m-saturated, from proposition 5 we obtain that there is an m-saturated I which is a model of $\text{PU}(M) \cup \{\phi\}$ and a lower bound of C with respect to \leq_M . But then, by lemma 9, there is a finite model F of ϕ such that $M \leftarrow F \leftarrow I$. Therefore, F is a lower bound of C and by applying Zorn's lemma we obtain stopperedness for the class of finite models. \square

We continue with a set of decidability results concerning the finite case. Firstly we prove that checking whether a specific model N is minimal with respect to a model M and ϕ , i.e. whether $N \in M *_f \phi$, is decidable. We will say that $M *_f \phi$ is non-trivial whenever $M \not\models \phi$ and $\text{PU}(M), \phi \not\vdash \perp$.

Lemma 10. *Let M, N be finite models and ϕ a sentence, such that $M *_f \phi$ is non-trivial. The decision problem of whether $N \in M *_f \phi$ is decidable.*

Proof. Since $M *_f \phi$ is non-trivial then if $M \leftarrow N$ then surely N is not minimal. So, we assume that $M \leftarrow N$. Now, N is minimal iff there is no other model $N' \in \text{mod}_f(\phi)$ such that $N' <_M N$. Assume N is not minimal. Then there exists $N' \in \text{mod}_f(\phi)$ such that $M \leftarrow N' \leftarrow N$ but $N' \not\rightarrow N$. By applying lemma 9 to the pair of models M, N' it follows that there exists a model $N'' \in \text{mod}_f(\phi)$ such that $M \leftarrow N'' \leftarrow N'$ and thus, $N'' \not\rightarrow N$. Therefore, N is not minimal iff there exists a model N'' of ϕ which is strictly smaller than N with respect to \leq_M and it has at most $f(M, \phi)$ states.

Consequently, given N, M and ϕ , we can enumerate all the models of ϕ that have at most $f(M, \phi)$ states, of which there is a finite number. For each model L we check the simulations $M \leftarrow L, L \leftarrow N, L \not\rightarrow N$. If we find a model that satisfies all those conditions, then N is not minimal. If we do not find one, then by using the result in the previous paragraph, N is minimal. \square

The next proposition characterises the structure of $M *_f \phi$ with respect to the ordering. It asserts that each equivalence class of models in $M *_f \phi$ with respect to \leq_M , contains a representative model of a bounded size.

Proposition 3. *Let M be a finite model and ϕ a formula such that $M *_f \phi$ is non-trivial. Then, there is a computable finite set of finite models $\Delta_{M,\phi} \subseteq M *_f \phi$ such that for any model $N \in M *_f \phi$ there is a model $N' \in \Delta_{M,\phi}$ such that $N \Leftrightarrow N'$ and $|N'| \leq f(M, \phi)$.*

Proof. Let $N \in M *_f \phi$. The application of lemma 9 gives us a model N' of ϕ such that $M \leftarrow N' \leftarrow N$ and $|N'| \leq f(M, \phi)$. But since N is minimal, it follows that $N \Leftrightarrow N'$. Thus $\Delta_{M,\phi}$ can be computed by enumerating the finite models of ϕ that have at most $f(M, \phi)$ states and checking them for minimality via lemma 10. \square

A corollary of the above is that given the premises of proposition 3, the number of equivalence classes of finite models that constitute $M *_f \phi$ is finite. We next examine the decidability of reasoning about the results of the operation.

Proposition 4. *Assume that $M *_f \phi$ is non-trivial. Let ψ be a formula in $\mathcal{L}_{\text{PU}} \cup \mathcal{L}_{\text{PE}}$. Then, the decision problem $M *_f \phi \models \psi$ is decidable.*

Proof. $M *_f \phi$ can be seen as the union of a (finite) set of equivalence classes of finite models under simulation. Let $E \subseteq M *_f \phi$ be such an equivalence class. For any two models $N_1, N_2 \in E$ it holds that $\text{PU}(N_1) = \text{PU}(N_2)$ and equivalently $\text{PE}(N_1) = \text{PE}(N_2)$. In other words, the problem of checking whether all models in $M *_f \phi$ satisfy ψ , where $\psi \in \mathcal{L}_{\text{PU}} \cup \mathcal{L}_{\text{PE}}$, reduces to checking whether for each equivalence class E in $M *_f \phi$, there is a model $N \in E$ such that $N \models \psi$. But $\Delta_{M,\phi}$ contains at least one model from each such equivalence class, so the problem is further reduced to whether $\Delta_{M,\phi} \models \psi$ or not. Since $\Delta_{M,\phi}$ is finite and computable, the problem is decidable. \square

Lastly, we mention some results that extend the ones in this section to serial models and the corresponding fragment to the PU formulas in the temporal logic $\text{CTL}^*[4]$, known as $\forall\text{CTL}^*[7]$ and its dual $\exists\text{CTL}^*$. Due to lack of space we omit the proofs that are in any case trivial extensions of the above. Note that in the following, the formula-argument of the operation $*$ remains a K_n formula.

- If M, N are serial models, then $M \leftarrow N$ implies $\forall\text{CTL}^*(M) \subseteq \forall\text{CTL}^*(N)$.
- If M, N are serial m-saturated models then $\forall\text{CTL}^*(M) \subseteq \forall\text{CTL}^*(N)$ implies $M \leftarrow N$.
- Lemma 5 extends to serial models too (the idea being that lemma 5 can be used as is but with T extended to $T \cup \Sigma$, where $\Sigma = \{ \Box^n \Diamond T \mid n \geq 0 \}$). It follows that proposition 1 extends to serial m-saturated models.
- Lemma 9 extends to serial models. The crucial point is that generated submodels of a serial model are serial as well. This implies that proposition 2 extends to serial finite models.

- Similarly, lemma 10 and proposition 3 extend to the case of serial finite models, by adding (decidable) checks for seriality in the models involved in the proofs.
- Proposition 4 can be extended to the following: Assume the conditions of proposition 3. Let ψ be a formula in $\forall\text{CTL}^* \cup \exists\text{CTL}^*$. Then, the decision problem $M *_f \phi \models \psi$ is decidable.

4 Conclusions and Further Work

Our results are positive and intuitive, showing that

- The refinement ordering \leq_M with respect to a model M is stoppered, and therefore has minimals, in the class of m -saturated models, for any set of sentences; and in the class of finite models, for any formula.
- The properties of minimal refinements over finite models are decidable.

A limitation of our framework is the expressiveness of the underlying logic K_n . Properties that involve e.g. transitivity, quantification over sets of states or computational paths in the model, cannot be expressed in K_n . To address this, we intend to extend our results to more expressive languages, and already have a preliminary set of results concerning K_n but with global validity in mind. Furthermore we intend to investigate the complexity of the algorithms we have presented.

References

- [1] Abadi, M. and L. Lamport, *The existence of refinement mappings*, Theoretical Computer Science **82** (1991), pp. 253–284.
- [2] Blackburn, P., M. de Rijke and Y. Venema, “Modal Logic,” Cambridge Tracts in Theoretical Computer Science **53**, Cambridge University Press, 2001.
- [3] Bouajjani, A., J. C. Fernandez, S. Graf, C. Rodriguez and J. Sifakis, *Safety for branching time semantics*, in: *Automata, Languages and Programming, 18th International Colloquium*, 1991, pp. 76–92.
- [4] Clarke, E. and E. Emerson, *Design and synthesis of synchronization skeletons using branching time temporal logic*, in: D. Kozen, editor, *Logics of Programs, Workshop*, Lecture Notes in Computer Science **131** (1981), pp. 52–71.
- [5] de Rijke, M., *Modal model theory*, Technical Report CS-R9517, CWI, Amsterdam (1995).

- [6] Gärdenfors, P., *Belief revision: An introduction*, in: P. Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 1992 pp. 1–28.
- [7] Grumberg, O. and D. E. Long, *Model checking and modular verification*, ACM Transactions on Programming Languages and Systems **16** (1994), pp. 843–871.
- [8] Hennessy, M. and R. Milner, *Algebraic laws for nondeterminism and concurrency*, Journal of the ACM **32** (1985), pp. 137–161.
- [9] Hollenberg, M., *Hennessy-Milner classes and process algebra*, in: Y. V. A. Ponse, M. de Rijke, editor, *Modal Logic and Process Algebra*, CSLI Publications, 1995 pp. 187–216.
- [10] Katsuno, H. and A. O. Mendelzon, *Propositional knowledge base revision and minimal change*, Artificial Intelligence **52** (1991), pp. 263–294.
- [11] Katsuno, H. and A. O. Mendelzon, *On the difference between updating a knowledge base and revising it*, in: P. Gärdenfors, editor, *Belief Revision*, Cambridge University Press, 1992 pp. 183–203.
- [12] Milner, R., *An algebraic definition of simulation between programs*, in: *Proceedings of the Second International Joint Conference on Artificial Intelligence* (1971), pp. 481–489.
- [13] Rosen, E., *Modal logic over finite structures*, Journal of Logic, Language and Information **6** (1997), pp. 427–439.