# Model Generation for Horn Logic with Stratified Negation 

Ethan K. Jackson and Wolfram Schulte<br>Microsoft Research, One Microsoft Way, Redmond, WA<br>\{ejackson, schulte\}@microsoft.com


#### Abstract

Model generation is an important formal technique for finding interesting instances of computationally hard problems. In this paper we study model generation over Horn logic under the closed world assumption extended with stratified negation. We provide a novel threestage algorithm that solves this problem: First, we reduce the relevant Horn clauses to a set of non-monotonic predicates. Second, we apply a fixed-point procedure to these predicates that reveals candidate solutions to the model generation problem. Third, we encode these candidates into a satisfiability problem that is evaluated with a state-of-the-art SMT solver. Our algorithm is implemented, and has been successfully applied to key problems arising in model-based design.


## 1 Introduction

Informally, model generation is a procedure that takes as input some mathematical statement $\psi$, and produces as output some data $M$ (a model) that, when substituted back into $\psi$, makes the statement true. For example, if $\psi$ is a boolean satisfiability problem, then $M$ is an assignment of boolean variables to truth values. Similarly, if $\psi$ is a set of linear inequalities, then $M$ is an assignment of variables to the real numbers. Note that model generation can be used to check satisfiability, but not all techniques for checking satisfiability are able to generate models. In this paper we study model generation for an important type of non-classical logic called Horn logic with stratified negation.

Horn logic has important applications in computer science and new applications continue to arise. Recently, Horn logic extended with negation as failure was used to formalize the non-context-free languages arising in modern software engineering methodologies[1] such as Model Driven Architecture[2][3] (MDA), Model Integrated Computing[4] (MIC), and Platform-based Design[5][6] (PBD). This particular application adds a new and interesting twist: An effective means of model generation is essential if the Horn paradigm is to be truly useful.

In this paper we present a novel approach to model generation for nonrecursive Horn logic extended with stratified negation. Our approach employs a three stage process:

1. We simplify the problem by reducing the relevant Horn clauses to a set of non-monotonic predicates that we call non-monotonic acceptors.
2. We apply a fixed-point procedure to the non-monotonic acceptors that reveals the candidate solutions to the model generation problem.
3. We encode these candidates into a satisfiability problem that is evaluated with the state-of-the-art SMT solver Z3.

We show that this procedure is sound, but incomplete.
This paper is organized into six sections. Section 2 informally describes the class of Horn logic targeted for model generation. Section 3 provides the key formal definitions. The first stage of the algorithm is explained in Section 4 as a modified form of backwards chaining. Section 5 describes the elimination of quantification over closed-worlds and the reduction to a boolean satisfiability problem. We conclude in Section 6.

## 2 Background and Running Example



Fig. 1. (Left) Scheduling abstraction (Right) Example scheduling instance

Model-based approaches to software engineering rely on high-levels of abstraction to simplify the design process. The left-hand side of Figure 1 shows a metamodel describing a simple abstraction layer. This diagram defines an abstract language for scheduling problems without regard to the particular details of the tasks being scheduled. This language contains objects of type Task and Processor. Tasks can be assigned to processors by directed edges of type TaskMap. Resource constraints between tasks prevent two tasks from being scheduled on the same processor. Resource constraints are modeled as undirected edges of type Constraint connecting tasks. The right-hand side of Figure 1 shows a member of this language. There are three tasks $T_{1}, T_{2}, T_{3}$ and two processors $P_{1}, P_{2}$. Tasks $T_{1}$ and $T_{2}$ have a resource constraint, as do tasks $T_{2}$ and $T_{3}$. Tasks $T_{1}, T_{3}$ are scheduled on processor $P_{1}$ while $T_{2}$ is scheduled on $P_{2}$.

Metamodels (and other artifacts) can expressed a set of axioms using Horn logic with stratified negation [7]. Consequently, model generation on this logic is a key tool for reasoning about abstraction layers. For example, we might demand a model generator to "Construct a model that contains three tasks and two processors". The procedure must find a model satisfying the rules of the abstraction that also meets this goal. Model generation is a difficult problem, as this example illustrates. A correct mapping from tasks to processors is precisely a graph
coloring problem[8] where the processors are the colors. This illustrates that any procedure capable of constructing non-trivial instances must solve difficult subproblems.

The task scheduling language is defined with the following set of non-recursive and stratified Horn clauses:

$$
\begin{align*}
\operatorname{task}(x) & \leftarrow \operatorname{taskmap}(x, y)  \tag{1}\\
\operatorname{processor}(y) & \leftarrow \operatorname{taskmap}(x, y)  \tag{2}\\
\operatorname{task}(x) & \leftarrow \operatorname{constraint}(x, y)  \tag{3}\\
\operatorname{task}(y) & \leftarrow \operatorname{constraint}(x, y)  \tag{4}\\
\operatorname{no\_ map}(\operatorname{task}(x)) & \leftarrow \operatorname{task}(x), \neg \operatorname{taskmap}(x, y)  \tag{5}\\
\text { bad_map }(\operatorname{task}(x), \operatorname{task}(y)) & \leftarrow \operatorname{taskmap}(x, z), \operatorname{taskmap}(y, z), \operatorname{constraint}(x, y) \tag{6}
\end{align*}
$$

The first four clauses declare that the end-points of task mappings and resource constraints always exist. Clause 5 deduces a term no_map $(\operatorname{task}(x))$ for any task $x$ that is not mapped to a processor. Clause 6 deduces bad_map $(\operatorname{task}(x), \operatorname{task}(y))$ anytime two tasks $x, y$ are improperly scheduled. We now examine the semantics of this logic in detail.

### 2.1 Classical Horn Logic

Classical Horn logic restricts first-order logic by requiring each conjunct of a DNF (disjunctive normal form) formula $\psi$ to have at most one non-negated literal. A collection of Horn formulas has a more natural representation in implicative normal form, as shown below:

$$
\begin{equation*}
\forall x, y, z \operatorname{taskmap}(x, z), \operatorname{taskmap}(y, z), \text { constraint }(x, y) \Rightarrow \operatorname{bad\_ sched}(x, y) \tag{7}
\end{equation*}
$$

This classical clause looks similar to Clause 6, however its meaning is quite different. For the sake of discussion, assume that $\operatorname{taskmap}(\cdot, \cdot)$, constraint $(\cdot, \cdot)$, bad_sched $(\cdot, \cdot)$ are predicates. From this clause we know that bad_sched $(x, y)$ must be true for tasks $x$ and $y$ that have resource constraint between them and are scheduled onto the same processor. Given two particular tasks $\mathrm{t}_{1}$, $\mathrm{t}_{2}$ without a resource constraint between them, what can we conclude about bad_sched $\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right)$ ? Rewriting Equation 7 yields:
$\forall z, \neg \operatorname{taskmap}\left(\mathrm{t}_{1}, z\right) \vee \neg \operatorname{taskmap}\left(\mathrm{t}_{2}, z\right) \vee \neg \operatorname{constraint}\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right) \vee$ bad_sched $\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right)$
Since $\neg \operatorname{constraint}\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right)$ is true, Equation 8 is satisfied without forcing a particular truth value to bad_sched $\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right)$. In other words, there exists a model satisfying Equation 7 for which tasks 1 and 2 are badly scheduled, but there also exists a model where they are not badly scheduled. Both of these possibilities exist because we have used classical implication. However, there exists a commonly employed extension to Horn logic that closes this loop-hole.

### 2.2 Closed World Assumption

The Closed World Assumption (CWA) is applied whenever a set of Horn clauses is intended to capture all the necessary information for a domain[9]. In order for CWA to work properly there must exist some information known to be true from the outset. These pieces of information are called facts, which are of the form true $\Rightarrow h$ where $h$ is a non-negated literal. For example:

$$
\begin{equation*}
\operatorname{true} \Rightarrow \operatorname{taskmap}\left(\mathrm{t}_{1}, \mathrm{p}_{1}\right) \tag{9}
\end{equation*}
$$

Intuitively, a predicate $f(x, y, z)$ is true for some $x, y, z$ if $f(x, y, z)$ is a fact or if there is a sequence of derivations starting at facts that force $f(x, y, z)$ to be true. If there is no such derivation, then $f(x, y, z)$ is false. This rule eliminates the case where tasks $t_{1}$ and $t_{2}$ are badly scheduled.

This slight adjustment to classical implication profoundly effects the underlying formal machinery by introducing a fixed-point operator $\widehat{\Gamma}$. This operator, called the immediate consequence operator, deduces new facts from existing facts using the clauses. All facts not deduced by $\widehat{\Gamma}$ are false, so any model generation procedure must reason over this operator. If $M$ is an initial set of facts and $\Lambda$ is a set of non-fact Horn clauses, then the set of facts deducible by $\Lambda$ is the least set $X$ such that $M \subseteq X$ and $X=\widehat{\Gamma}(X)$. CWA can also be understood from a different angle under the name existential fixed-point logic [10].

The Closed World Assumption is used in most applications of Horn logic, so it must be taken into account by any model generation procedure. However, CWA forces a rephrasing of the model generation problem: Let $\Lambda$ be a set of non-fact Horn clauses and a goal $G=\left\{l_{1}, l_{2}, \ldots, l_{n}\right\}$ be a set of non-negated literals. Loosely speaking, the finite model generation problem is to find a finite set of facts $M$ so that (1) $M \subseteq X,(2) X$ is a least fixed-point of $\widehat{\Gamma}$, and (3) $X$ contains the goal literals (with respect to some substitution.) Thus, any model generation procedure must reason carefully about the fixed-points of $\widehat{\Gamma}$.

### 2.3 Negation-as-Failure

Classical Horn logic (without CWA) restricts the use of negation, which restricts the expressiveness of the fragment. Negation-as-failure (NAF) attempts to reintroduce a form of negation that is compatible with CWA and does not recreate full first-order logic. However, this new form of negation is very different from its classical counterpart. Intuitively, a negated literal $\neg l$ is true if $l$ cannot be proved true under Horn logic with CWA. Thus, negation is defined in terms of a proof procedure.

In order to distinguish our Horn clauses from the classical fragment we write a clause this way:

$$
\begin{equation*}
h \leftarrow s_{1}, s_{2}, \ldots, s_{m}, \quad \neg t_{1}, \neg t_{2}, \ldots, \neg t_{k} \tag{10}
\end{equation*}
$$

The literal $h$ is called the head of the clause and $\left\{s_{1}, \ldots, s_{m}, t_{1}, \ldots, t_{k}\right\}$ is the tail of the clause. Each $\neg t_{i}$ is a negated literal where negation refers to nonclassical NAF. Consider Clause 5 containing the negated term $\neg \operatorname{taskmap}(x, y)$.

This negation does not directly ask if $\operatorname{taskmap}(x, y)$ is false for some $x, y$. Instead, it asks if $\operatorname{taskmap}(x, y) \notin \widehat{\Gamma}(M)$ for some $x, y$. Unlike classical negation, NAF must be used carefully otherwise logical inconsistencies can arise. For example, under NAF we might simultaneously conclude $f \in \widehat{\Gamma}(M) \wedge f \notin \widehat{\Gamma}(M)$ for some fact $f$. This is a more dangerous inconsistency than $b \wedge \neg b$ for some boolean variable $b$, which has a well-defined meaning. Much work has been done on generalized forms of NAF that do not suffer from inconsistencies [11][7][12]. We avoid these problems by using only a restricted form of NAF called stratified negation. Our approach to handling NAF is similar to the non-monotonic rules described in [13].

## 3 Definitions

We now formally describe the style of Horn logic for which we generate models; this logic incorporates both CWA and NAF. Note that our definitions are biased to make the presentation of model generation simpler.

### 3.1 Basic Concepts

Let $\Upsilon$ denote a finite signature, $\Sigma$ an infinite alphabet of constants, and $\mathcal{V}$ a infinite alphabet of variable names. We use the letters $f, g, h$ for variables ranging over function symbols of some signature $\Upsilon$. We use typewriter script to denote constants from $\Sigma$. Finally, we use $x, y, z$ for variables ranging over terms. Let $\operatorname{arity}(f)$ denote the arity of some function symbol $f$. A term is a combination of function symbols, constants, and variables:

Definition 1. Given $\Upsilon, \Sigma$, and $\mathcal{V}$, the set of all finite terms $\mathcal{T}$ is defined inductively

1. Each $c \in \Sigma$ is a term
2. Each $x \in \mathcal{V}$ is a term
3. If $f \in \Upsilon$ and $t_{1}, t_{2}, \ldots, t_{\text {arity }(f)} \in \mathcal{T}$ then $f\left(t_{1}, t_{2}, \ldots, t_{\text {arity }(f)}\right)$ is a term.

If it is unclear from context, we write $\mathcal{T}(\Upsilon)$ to denote the finite terms constructed from function symbols of signature $\Upsilon$. A ground term is a term without variables; we use $\mathcal{T}_{G}$ to denote the set of all ground terms. If $t$ is a term, then $s$ is a subterm of $t$ (written $s \sqsubseteq t$ ) if $s=t$ or $s$ is a subterm of one of the arguments of $t$. If $t$ is a term, then $\operatorname{vars}(t)$ is the set of subterms that are also variable names. Similarly, consts $(t)$ is the set of subterms that are also constants. These functions are extended to sets of terms $S \subseteq \mathcal{T}$ in the natural way: $\operatorname{vars}(S)$ (or consts $(S)$ ) are the variables (or constants) appearing in a set of terms.

### 3.2 Substitutions and Unifiers

Terms are related to one another through special homomorphisms called substitutions.

Definition 2. A substitution $\varphi: \mathcal{T} \rightarrow \mathcal{T}$ is a mapping from terms to terms such that:

1. $\varphi$ fixes constants, i.e. $\forall c \in \Sigma, \varphi(c)=c$.
2. $\varphi$ is a term homomorphism, i.e. $\varphi\left(f\left(t_{1}, t_{2}, \ldots, t_{n}\right)\right)=f\left(\varphi\left(t_{1}\right), \varphi\left(t_{2}\right), \ldots, \varphi\left(t_{n}\right)\right)$.

Let $\Phi$ be the set of all substitutions for some $\mathcal{T}$. Two terms $s, t$ are said to unify if there exists a substitution $\varphi$ that makes them the same: $\varphi(s)=\varphi(t)$. Essentially, a substitution makes two terms the same by replacing variables in the terms with new subterms. The essence of this replacement is easily characterized in terms of the kernel of $\varphi$.
Definition 3. The kernel of a homomorphism $\varphi$ (e.g. a substitution) is:

$$
\begin{equation*}
\operatorname{ker} \varphi=\left\{(s, t) \in \mathcal{T}^{2} \mid \varphi(s)=\varphi(t)\right\} \tag{11}
\end{equation*}
$$

The kernel characterizes which subterms are equated by a substitution without regard to the particular values assigned to variables. Some important properties of kernels are: (1) Every kernel is an equivalence relation. (2) The intersection of two equivalence relations is also an equivalence relation. (3) The least equivalence relation $\Theta$ containing two equivalence relations $\Theta_{1}, \Theta_{2}$ is the transitive closure of $\left(\Theta_{1} \cup \Theta_{2}\right)$. This shall be written $\Theta=\Theta_{1} \oplus \Theta_{2}$.

The most general unifiers (mgu) of two terms $s, t$ is an equivalence relation between the variables of $s$ and $t$ that must hold for any substitution unifying $s$ and $t$. This equivalence relation represents the weakest set of constraints over the variables of $s$ and $t$ that ensures unification. The most general unifiers have the following properties:

Lemma 1. Given two terms $s, t$ that unify, let mgu( $s, t)$ denote the most general unifiers.

1. The mgu( $s, t)$ is unique.
2. $m g u(s, t)=\bigcap_{\varphi}\{\operatorname{ker} \varphi \mid \varphi(s)=\varphi(t)\}$

If terms $s$ and $t$ do not unify, then we write $m g u(s, t)=\emptyset$.

### 3.3 Horn Logic with CWA and NAF

Given $\Upsilon, \Sigma, \mathcal{V}$ a Horn clause $\lambda$ is a triple $\lambda=(h, P, N)$ where $h$ is a term and $P, N$ are sets of terms. $P=\left\{s_{1}, \ldots, s_{m}\right\}$ is the set of non-negated tail terms and $N=\left\{t_{1}, \ldots, t_{k}\right\}$ is the set of negated tail terms. A Horn clause is written:

$$
\begin{equation*}
h \leftarrow s_{1}, s_{2}, \ldots, s_{m}, \quad \neg t_{1}, \neg t_{2}, \ldots, \neg t_{k} \tag{12}
\end{equation*}
$$

Furthermore, $P$ must be non-empty and $\operatorname{vars}(h) \subseteq \operatorname{var} s(P)$. (We shall explain these restrictions shortly.) Let $\Lambda$ be a finite set of Horn clauses, then there exists a binary relation $\prec$ over clauses, where $\left(h^{\prime}, P^{\prime}, N^{\prime}\right) \prec(h, P, N)$ if there exists some $s_{i} \in P$ or $t_{j} \in N$ that unifies with $h^{\prime}$.

Definition 4. Let $\Lambda$ be a finite set of clauses, then $\Lambda$ is non-recursive and stratified $i f \prec$ is a strict partial order.

Restricting $\prec$ to a strict partial order yields a simple semantics for evaluating the set of facts derivable by $\Lambda$. Order the clauses $\lambda_{1}, \lambda_{2}, \ldots, \lambda_{k}$ so that $\lambda_{i} \prec \lambda_{j}$ implies $i<j$, then for each clause define an immediate consequence operator:

$$
\begin{equation*}
\widehat{\Gamma}_{i}(X)=\bigcup_{\varphi}\left\{\varphi\left(h_{i}\right) \mid \alpha\left(P_{i}, N_{i}, \varphi, X\right)\right\} \cup X \tag{13}
\end{equation*}
$$

This equation states that the facts deducible by a single clause $\lambda_{i}$ are calculated by finding all the substitutions that satisfy a special predicate $\alpha_{i}\left(P_{i}, N_{i}, \varphi, X\right)$; each substitution is applied to the head $h_{i}$ to derive a new fact. Earlier we restricted the variables of $h_{i}$ to be a subset of the variables of $P_{i}$, so each substitution maps $h_{i}$ to a well-defined ground term.

The predicate $\alpha\left(P_{i}, N_{i}, \varphi, X\right)$ captures the CWA and NAF semantics.
Definition 5. $\alpha: \mathcal{P}(\mathcal{T})^{2} \times \Phi \times \mathcal{P}(\mathcal{T}) \rightarrow \mathbb{B}$ is called $a$ non-monotonic acceptor:

$$
\begin{equation*}
\alpha(P, N, \varphi, X) \stackrel{\text { def }}{=}(\varphi(P) \subseteq X) \wedge \forall \varphi^{\prime}\left[\left(\varphi(P)=\varphi^{\prime}(P)\right) \Rightarrow\left(\varphi^{\prime}(N) \cap X=\emptyset\right)\right] \tag{14}
\end{equation*}
$$

The acceptor $\alpha(P, N, \varphi, X)$ is true for some substitution $\varphi$ and some set of terms (e.g. facts) $X$ if the positive terms $P$ can be found in the set of facts through the substitution $\varphi$. The negative terms $N$ must not be found in the facts $X$ for any extension of $\varphi$ to $\varphi^{\prime}$ that agrees on $P$.

Lemma 2. Let $\Lambda$ be a finite set of non-recursive and stratified Horn clauses, and $M$ a finite set of ground terms. Let the clauses of $\Lambda$ be ordered $\lambda_{1}, \lambda_{2}, \ldots, \lambda_{k}$ to respect $\prec$ then:

1. The set of all facts deducible from $M$ by $\Lambda$ is $\Gamma(M)$ where:

$$
\begin{equation*}
\Gamma(M)=\widehat{\Gamma}_{k}\left(\ldots \widehat{\Gamma}_{2}\left(\widehat{\Gamma}_{1}(M)\right) \ldots\right) \tag{15}
\end{equation*}
$$

2. It can be decided in finite time if any ground term $t_{g} \in \Gamma(M)$, i.e. the logic is decidable.

### 3.4 The Model Generation Problem

Solving the model generation problem requires the construction of a set of facts $M$ that satisfies a goal. A goal $G=\left(P_{G}, N_{G}\right)$ is comprised of two sets of terms: the positive terms $P_{G}$, and the negative terms $N_{G}$. A goal is satisfied if there exists some $M$ such that all the facts deduced from $M$ (i.e., $\Gamma(M)$ ) include $P_{G}$ and do not include $N_{G}$. More precisely, $M$ satisfies the goal if $\exists \varphi, \alpha\left(P_{G}, N_{G}, \varphi, \Gamma(M)\right)$ holds.

In order to construct meaningful solutions, the model generation procedure must know which terms are allowed to appear as facts. Consider the problem of creating a badly scheduled set of tasks for the scheduling abstraction:

$$
G=\left(\left\{b a d \_m a p(x, y)\right\}, \emptyset\right)
$$

then we expect the model generation procedure return a solution similar to this one:

$$
M=\left\{\begin{array}{c}
\operatorname{task}\left(\mathrm{t}_{1}\right), \operatorname{task}\left(\mathrm{t}_{2}\right), \operatorname{processor}\left(\mathrm{p}_{1}\right), \\
\operatorname{taskmap}\left(\mathrm{t}_{1}, \mathrm{p}_{1}\right), \operatorname{taskmap}\left(\mathrm{t}_{2}, \mathrm{p}_{1}\right), \operatorname{constraint}\left(\mathrm{t}_{1}, \mathrm{t}_{2}\right)
\end{array}\right\} .
$$

Without additional information, the solution $M=\left\{\operatorname{bad\_ map}(\mathrm{c}, \mathrm{c})\right\}$ is also trivially valid. This extra information is expressed by partitioning the signature $\Upsilon$ into two parts: the fact signature $\Upsilon_{F}$ and the derived signature $\Upsilon_{D}$. We call a term $t \in \mathcal{T}\left(\Upsilon_{F}\right)$ a fact term and call all other terms derived terms. The model generation procedure only considers solutions that are sets of fact terms. For example, the partitioning:

$$
\begin{aligned}
& \Upsilon_{F}=\{\operatorname{task}(\cdot), \text { processor }(\cdot), \operatorname{taskmap}(\cdot, \cdot), \text { constraint }(\cdot, \cdot)\}, \\
& \Upsilon_{D}=\{\text { no_map }(\cdot), \text {,bad_map }(\cdot, \cdot)\}
\end{aligned}
$$

forces all solutions to be built from tasks and processors. The goal $G$ contains a derived term $\operatorname{bad\_ map}(x, y)$, but the solution $M$ will never contain this term directly. We now define the finite model generation problem.

Definition 6. The finite model generation problem - Given:

1. A finite signature $\Upsilon$ partitioned into $\Upsilon_{F} \neq \emptyset$ and $\Upsilon_{D}$,
2. A finite set of clauses $\Lambda$ that are non-recursive and stratified,
3. A goal $G=\left(P_{G}, N_{G}\right)$ where $P_{G}, N_{G}$ are finite subsets of terms.

Construct a finite set of ground terms $M \subset \mathcal{T}_{G}\left(\Upsilon_{F}\right)$ so that

$$
\begin{equation*}
\exists \varphi \alpha\left(P_{G}, N_{G}, \varphi, \Gamma(M)\right) \tag{16}
\end{equation*}
$$

## 4 Utilizing Backwards Chaining

The semantics of non-recursive and stratified Horn logic has a succinct characterization in terms of $\Gamma$. However, it is difficult to construct a set of constraints from $\Gamma$ that guide model generation. Fortunately, there is a well-known technique for computing the truth values of goals, called backwards chaining, which addresses this problem. Imagine that a set of facts $M$ is already known, and the only task is to check if a goal $G$ is satisfied by these facts. Backwards chaining works backwards from the goal, through the clauses, to the facts $M$ to check satisfiability. If the goal is satisfied, then the procedure yields a proof tree showing exactly how the facts derive the goal. Formally, this process is called $S L D$ resolution [14] for Horn logic and SLDNF resolution [15] for Horn logic with NAF. These resolution procedures are sound and complete for non-recursive and stratified Horn logic. We modify SLDNF resolution to return "possible" proof trees, and then search for models that satisfy these proof trees. We utilize soundness/completeness results [16] to argue soundness for model generation.

The key modification to SLDNF is a new termination condition that does not rely on $M$. Typical backwards chaining terminates when it encounters a fact, i.e. a clause of the form $f \leftarrow t r u e$. This termination condition must be modified
for model generation because initially no facts are known. We modify backwards chaining so that it terminates when a fact term is encountered, even though this fact term may not exist in the solution $M$. Let a clause $\lambda \in \Lambda$ be partitioned as follows:

$$
\begin{equation*}
h \leftarrow p_{1}, \ldots, p_{m}, \quad u_{1}, \ldots, u_{m^{\prime}}, \quad \neg n_{1}, \ldots, \neg n_{k}, \quad \neg w_{1}, \ldots, \neg w_{k^{\prime}} \tag{17}
\end{equation*}
$$

where (1) each $p_{i}$ is a positive fact term, (2) each $u_{i}$ is a positive derived term, (3) each $n_{i}$ is a negative fact term, and (4) each $w_{i}$ is a negative derived term.

Associated with each clause $\lambda$ is a backwards chaining predicate $\beta_{\lambda}(\varphi, M, \Theta)$ where $M$ is a set of terms, $\varphi$ is a substitution, and $\Theta$ is an equivalence relation.

Definition 7. Let $\lambda$ be a clause, then associate with $\lambda$ a backwards chaining predicate ${ }^{1} \beta_{\lambda}(\varphi, M, \Theta)$ :
$\beta_{\lambda}(\varphi, M, \Theta) \stackrel{\text { def }}{=}$

1. $(\Theta \subseteq \operatorname{ker} \varphi) \wedge$
2. $\alpha\left(\left\{p_{1}, p_{2}, \ldots, p_{k}\right\},\left\{n_{1}, n_{2}, \ldots, n_{m}\right\}, \varphi, M\right) \wedge$
3. $\bigwedge_{1 \leq i \leq m^{\prime}}\left(\bigvee_{m g u\left(u_{i}, h_{\lambda^{\prime}}\right) \neq \emptyset} \beta_{\lambda^{\prime}}\left(\varphi, M, \Theta \oplus m g u\left(u_{i}, h_{\lambda^{\prime}}\right)\right)\right) \wedge$
4. $\forall \varphi^{\prime}\left[\begin{array}{l}\left(\forall v \in \operatorname{vars}\left(P_{\lambda}\right) \varphi^{\prime}(v)=\varphi(v)\right) \Rightarrow \\ \left(\bigwedge_{1 \leq j \leq k^{\prime} m g u\left(w_{j}, h_{\lambda^{\prime \prime}}\right) \neq \emptyset} \neg \beta_{\lambda^{\prime \prime}}\left(\varphi^{\prime}, M, \Theta \oplus \operatorname{mgu}\left(w_{j}, h_{\lambda^{\prime \prime}}\right)\right)\right)\end{array}\right]$

The backwards chaining predicate is defined recursively and terminates on fact terms; these parts of the tail simplify to non-monotonic acceptors (Def. 7.2 ). On the other hand, derived terms must be understood through additional clauses. The backwards chaining process recurses into derived terms by locating clauses that unify with these terms. The equivalence relation $\Theta$ is used to collect unification constraints during this process. For every positive derived $u_{i}$ there must exist some unifying clause $\lambda^{\prime}$ so that $\beta_{\lambda^{\prime}}$ is satisfied (Def. 7.3). Contrarily, every negative derived term $w_{j}$ must have no clause $\lambda^{\prime \prime}$ that derives $w_{j}$ for any extension of $\varphi$ to $\varphi^{\prime}$ agreeing on positive variables (Def. 7.4). It is possible that some unification constraints cannot be satisfied by any substitution $\varphi$. If this occurs then Def. 7.1 fails to hold. We assume that each time a clause $\lambda^{\prime}$ (or $\lambda^{\prime \prime}$ ) is examined for unification its variables are renamed to new variables that have not appeared before. This is called standardizing apart, and it prevents clauses from improperly interacting through variable names.

Backwards chaining is used to reduce any goal into a set of non-monotonic acceptors. However, there is one problem with the simple definition presented

[^0]here: It does not recurse through clauses with heads that are fact terms. (Consider clause 1 from the previous example.) This definition assumes that fact terms do not appear as heads. Fortunately, we can always rewrite the clauses of $\Lambda$ to enforce this rule; this is discussed later. For the moment, assume that facts do not appear as heads then the following important theorem holds:

Theorem 1. Given $\Upsilon_{F}, \Upsilon_{R}, \Lambda$ such that fact terms do not appear as heads, and a goal $G$, then

$$
\begin{equation*}
\forall M \subset \mathcal{T}\left(\Upsilon_{F}\right), \forall \varphi \quad \alpha\left(P_{G}, N_{G}, \varphi, \Gamma(M)\right) \Leftrightarrow \beta_{G}\left(\varphi, M, \mathcal{I D}_{\mathcal{T}}\right) \tag{18}
\end{equation*}
$$

for $M$ finite and $\Lambda$ non-recursive and stratified.
This theorem shows that evaluating the non-monotonic acceptor over all the facts deducible from $\Gamma$ gives the same result as working backwards from the goal $G$ through the backwards chaining predicates. We use this result to eliminate the fixed-point operator $\Gamma$ from the model generation problem. Note that the backwards chaining process is initiated without any constraints on the variables as described by the identity relation $\mathcal{I D}_{\mathcal{T}}=\{(t, t) \mid t \in \mathcal{T}\}$.

### 4.1 Simplification of Backwards Chaining

The backwards chaining formulation eliminates $\Gamma$, but generates many constraints across many recursions. In this section we show how to aggregate these constraints into convenient pieces. To facilitate this discussion we give names to particular parts of the backwards chaining predicate:

$$
\begin{align*}
\omega\left(\varphi, \varphi^{\prime}, P\right) & \stackrel{\text { def }}{=}\left(\forall v \in \operatorname{vars}(P) \varphi^{\prime}(v)=\varphi(v)\right)  \tag{19}\\
\psi_{\lambda}^{-}\left(\varphi^{\prime}, M, \Theta\right) & \stackrel{\text { def }}{=} \bigwedge_{1 \leq j \leq k^{\prime}} \bigwedge_{m g u\left(w_{j}, h_{\lambda^{\prime \prime}}\right) \neq \emptyset} \neg \beta_{\lambda^{\prime \prime}}\left(\varphi^{\prime}, M, \Theta \oplus m g u\left(w_{j}, h_{\lambda^{\prime \prime}}\right)\right) \tag{20}
\end{align*}
$$

Definition 7.4 becomes $\forall \varphi^{\prime} \omega\left(\varphi, \varphi^{\prime}, P_{\lambda}\right) \Rightarrow \psi_{\lambda}^{-}\left(\varphi^{\prime}, M, \Theta\right)$.
Consider the action of the goal backwards chaining predicate $\beta_{G}(\varphi, M, \Theta)$, as shown in Figure 2. The predicate $\beta_{G}$ introduces a non-monotonic acceptor $\alpha_{1}$ and some constraints on the kernel of $\varphi$ via $\Theta_{1}$. (Note that we index the


Fig. 2. A single expansion of the backwards chaining predicate for some unification choices of positive derived terms.
constraints $1,2, \ldots$ as the recursion proceeds.) Similarly, a subformula containing $\psi_{1}^{-}$is introduced due to negative derived terms. The positive derived terms $u_{1}, u_{2}, \ldots, u_{m^{\prime}}$ act as choice-points, because there may exist many clauses that unify with each $u_{i}$. Consider some choice of unifications for each $u_{i}$, then the recursion introduces more kernel constraints, non-monotonic acceptors, and negative subformulas. Let $\widehat{\beta}$ be an expansion of some $\beta$ for particular unification choices of the positive derived terms appearing through the recursion. Then this expansion has the following form:

$$
\begin{align*}
& \widehat{\beta}(\varphi, M, \Theta)= \\
& \left(\bigwedge_{i}\left(\Theta_{i} \subseteq \operatorname{ker} \varphi\right)\right) \wedge\left(\bigwedge_{i} \alpha\left(P_{i}, N_{i}, \varphi, M\right)\right) \wedge\left(\bigwedge_{i}^{\forall \varphi^{\prime} \omega\left(\varphi, \varphi^{\prime}, V_{i}^{+}\right) \Rightarrow} \begin{array}{l}
\psi_{i}^{-}\left(\varphi^{\prime}, M, \Theta_{i}\right)
\end{array}\right) \tag{21}
\end{align*}
$$

The following lemmas help to simplify the expansion.
Lemma 3. Non-monotonic acceptors compose over conjunction for fixed $\varphi$ and $X$.

$$
\begin{equation*}
\alpha\left(P_{a}, N_{a}, \varphi, X\right) \wedge \alpha\left(P_{b}, N_{b}, \varphi, X\right)=\alpha\left(P_{a} \cup P_{b}, N_{a} \cup N_{b}, \varphi, X\right) \tag{22}
\end{equation*}
$$

Lemma 4. Constraints on the kernel of $\varphi$ compose over conjunction.

$$
\begin{equation*}
\left(\Theta_{a} \subseteq \operatorname{ker} \varphi\right) \wedge\left(\Theta_{b} \subseteq \operatorname{ker} \varphi\right)=\left(\Theta_{a} \oplus \Theta_{b}\right) \subseteq \operatorname{ker} \varphi \tag{23}
\end{equation*}
$$

Applying the lemmas simplifies Equation 21 to:

$$
\begin{equation*}
\widehat{\beta}(\varphi, M, \Theta)=\left(\Theta^{\prime} \subseteq \operatorname{ker} \varphi\right) \wedge \alpha\left(P^{\prime}, N^{\prime}, \varphi, M\right) \wedge\binom{\forall \varphi_{i}^{\prime} \omega\left(\varphi, \varphi^{\prime}, V_{i}^{+}\right) \Rightarrow}{\psi_{i}^{-}\left(\varphi^{\prime}, M, \Theta_{i}\right)} \tag{24}
\end{equation*}
$$

where

$$
\begin{array}{r}
\Theta^{\prime}=\Theta \oplus\left(\bigoplus \Theta_{i}\right) \\
P^{\prime}=\bigcup P_{i}, \quad N^{\prime}=\bigcup N_{i} \tag{26}
\end{array}
$$

In summary, for a particular set of unification choices the backwards chaining reduces to:

1. Constraints on the kernel of $\varphi$, which equate variables,
2. A single non-monotonic acceptor containing only fact terms,
3. A number of backwards chaining predicates for negative derived terms.

A clause $\lambda$ may have an exponential number of expansions $\widehat{\beta}$. Label these expansions $\widehat{\beta}_{1}, \ldots, \widehat{\beta}_{c_{\lambda}}$, then they relate to the original predicate through disjunction:

$$
\begin{equation*}
\beta_{\lambda}(\varphi, M, \Theta)=\bigvee_{1 \leq i \leq c_{\lambda}} \widehat{\beta}_{i}(\varphi, M, \Theta) \tag{27}
\end{equation*}
$$

This decomposition also allows the $\psi^{-}$terms to be rewritten in terms of the expansion:

$$
\psi^{-}(\varphi, M, \Theta)=\bigwedge_{\operatorname{mgu}\left(w_{i}, h_{\lambda^{\prime}}\right) \neq \emptyset} \bigwedge_{1 \leq j \leq c_{\lambda^{\prime}}} \forall \varphi^{\prime}\left[\begin{array}{l}
\omega\left(\varphi, \varphi^{\prime}, V^{+}\right) \Rightarrow  \tag{28}\\
\neg \widehat{\beta}_{j}\left(\varphi^{\prime}, M, \Theta \oplus \operatorname{mgu}\left(w_{i}, h_{\lambda^{\prime}}\right)\right)
\end{array}\right]
$$

The decomposition of $\beta_{\lambda}$ shows that each $\psi^{-}$term will expand into some number of non-monotonic acceptors depending on the depth of the negation, which is certainly finite. Unlike the positive derived terms, each $\psi^{-}$must examine all the relevant $\widehat{\beta}$ predicates to ensure that the negated derived term is not satisfied. Furthermore, the simplification lemmas cannot be directly applied to expansions of $\psi^{-}$because the non-monotonic acceptors appear in negated form. In the following sections we use these expansions to generate models from the backwards chaining proof trees.

### 4.2 Restratification

The previous analysis assumed that backwards chaining terminates at fact terms. This assumption can be violated if $\Lambda$ contains clauses with fact terms as heads. The clause $\operatorname{task}(x) \leftarrow \operatorname{taskmap}(x, y)$ is an example. Fortunately, there is a simple syntactic operation that soundly manipulates $\Lambda$ so that no fact terms appear as heads. We call this process restratification, because it changes the ordering $\prec$.

Definition 8. Given $\Upsilon_{F}, \Upsilon_{D}, \Lambda$, and a goal $G$, then the restratified system is $\Upsilon_{F}, \Upsilon_{D}^{*}, \Lambda^{*}, G^{*}$ where:

1. Introduce a new unary derived function symbol restrat $(\cdot)$ to $\Upsilon_{R}$.

$$
\begin{equation*}
\Upsilon_{D}^{*}=\Upsilon_{D} \cup\{\operatorname{restrat}(\cdot)\} \tag{29}
\end{equation*}
$$

2. For each clause $\lambda \in \Lambda$ where the head $h$ is a fact term, add the modified clause $\lambda^{*}$ to $\Lambda^{*}$ :

$$
\begin{equation*}
\operatorname{restrat}(h) \leftarrow \neg h, \quad s_{1}, s_{2}, \ldots, s_{n}, \quad \neg t_{1}, \neg t_{2}, \ldots, \neg t_{m} \tag{30}
\end{equation*}
$$

3. For each clause $\lambda \in \Lambda$ where the head $h$ is derived term, add $\lambda$ to $\Lambda^{*}$
4. Modify the goal $G=\left(P_{G}, N_{G}\right)$ to include the negative derived term $\neg$ restrat $(x)$ where $x$ is a variable that does not appear in $G$.

$$
\begin{equation*}
G^{*}=\left(P_{G}, N_{G} \cup\{\operatorname{restrat}(x)\}\right) \tag{31}
\end{equation*}
$$

Lemma 5. If $\Lambda$ is a set of non-recursive and stratified Horn clauses, then the restratified clauses $\Lambda^{*}$ are also non-recursive and stratified where no clause has a fact term as head.

Theorem 2. The models that satisfy $G$ are related to the models that satisfy $G^{*}$ according to:

$$
\begin{align*}
& \forall M \subset \mathcal{T}_{G}\left(\Upsilon_{F}\right),\left(\Gamma_{\Lambda}(M) \cap \mathcal{T}_{G}\left(\Upsilon_{F}\right)=M\right) \Rightarrow \\
& \left(\exists \varphi \alpha\left(P_{G}, N_{G}, \varphi, \Gamma_{\Lambda}(M)\right) \Leftrightarrow \exists \varphi^{\prime} \alpha\left(P_{G}^{*}, N_{G}^{*}, \varphi^{\prime}, \Gamma_{\Lambda^{*}}(M)\right)\right) \tag{32}
\end{align*}
$$

If $M$ is a set of ground fact terms such that $\Gamma_{\Lambda}(M)$ does not grow the number of fact terms, then the restratified system will be in agreement with $G$. Conversely, when models are found that satisfy the restratified system, then these models do not grow fact terms under the original system. Of course, ground derived terms can still grow under either system. This is not a limitation, because solutions to $G$ that are not solutions to $G^{*}$ will accumulate fact terms under $\Gamma_{\Lambda}$ until the set of ground fact terms is exactly a solution to $G^{*}$.

### 4.3 Generating Schedules: Part 1

We now apply these techniques to the running example of a scheduling abstraction. Our goal is to find a model that contains three tasks $T_{1}, T_{2}, T_{3}$ and two processors $P_{1}, P_{2}$ so that the tasks are scheduled onto the processors. Furthermore tasks $T_{1}$ and $T_{2}$ cannot be on the same processor; the same constraint holds for task $T_{2}$ and $T_{3}$. In order to make the example more interesting the tasks are not introduced in the goal:

$$
\begin{align*}
& P_{G}=\left\{\operatorname{processor}\left(p_{1}\right), \text { processor }\left(p_{2}\right), \text { constraint }\left(t_{1}, t_{2}\right), \text { constraint }\left(t_{2}, t_{3}\right)\right\} \\
& P_{N}=\left\{\operatorname{no\_ map}(x), \operatorname{bad\_ map}(y, z)\right\} \tag{33}
\end{align*}
$$

for $\Upsilon_{R}=\left\{\right.$ no_map $\left.(\cdot), \operatorname{bad\_ map}(\cdot, \cdot)\right\}$ and all other function symbols in $\Upsilon_{F}$.
Applying restratification to the original clauses yields:

$$
\begin{align*}
\operatorname{restrat}(\operatorname{task}(x)) & \leftarrow \operatorname{taskmap}(x, y), \neg \operatorname{task}(x)  \tag{34}\\
\operatorname{restrat}(\operatorname{processor}(y)) & \leftarrow \operatorname{taskmap}(x, y), \neg \operatorname{processor}(y)  \tag{35}\\
\operatorname{restrat}(\operatorname{task}(x)) & \leftarrow \operatorname{constraint}(x, y), \neg \operatorname{task}(x)  \tag{36}\\
\operatorname{restrat}(\operatorname{task}(y)) & \leftarrow \operatorname{constraint}(x, y), \neg \operatorname{task}(y)  \tag{37}\\
\operatorname{no\_ map}(\operatorname{task}(x)) & \leftarrow \operatorname{task}(x), \neg \operatorname{taskmap}(x, y)  \tag{38}\\
\operatorname{bad\_ map}(\operatorname{task}(x), \operatorname{task}(y)) & \leftarrow \operatorname{taskmap}(x, z), \operatorname{taskmap}(y, z), \operatorname{constraint}(x, y) \tag{39}
\end{align*}
$$

and the restratified goal is $G^{*}=\left(P_{G}, N_{G} \cup\{\operatorname{restrat}(w)\}\right)$.
Next, the backwards chaining predicates are expanded until the goal is expressed by a system of non-monotonic acceptors defined with only fact terms.

In this example the $\omega$ terms are trivial and have been removed.

$$
\begin{align*}
\exists M, & \exists \varphi \beta_{G^{*}}\left(\varphi, M, \mathcal{I} \mathcal{D}_{\mathcal{T}}\right)= \\
& \exists M, \exists \varphi, \forall \varphi^{\prime} \alpha\left(P_{G}, \emptyset, \varphi, M\right) \wedge \\
& \neg \alpha\left(\left\{\operatorname{task}\left(x_{1}\right)\right\},\left\{\operatorname{taskmap}\left(x_{1}, y_{1}\right)\right\}, \varphi^{\prime}, M\right) \wedge \\
& \neg \alpha\left(\left\{\operatorname{taskmap}\left(x_{2}, y_{2}\right)\right\},\left\{\operatorname{task}\left(x_{2}\right)\right\}, \varphi^{\prime}, M\right) \wedge \\
& \neg \alpha\left(\left\{\operatorname{taskmap}\left(x_{3}, y_{3}\right)\right\},\left\{\operatorname{processor}\left(y_{3}\right)\right\}, \varphi^{\prime}, M\right) \wedge  \tag{40}\\
& \neg \alpha\left(\left\{\operatorname{constraint}\left(x_{4}, y_{4}\right)\right\},\left\{\operatorname{task}\left(x_{4}\right)\right\}, \varphi^{\prime}, M\right) \wedge \\
& \neg \alpha\left(\left\{\operatorname{constraint}\left(x_{5}, y_{5}\right)\right\},\left\{\operatorname{task}\left(y_{5}\right)\right\}, \varphi^{\prime}, M\right) \wedge \\
& \neg \alpha\left(\left\{\operatorname{taskmap}\left(x_{2}, z_{2}\right), \operatorname{taskmap}\left(y_{2}, z_{2}\right), \operatorname{constraint}\left(x_{2}, y_{2}\right)\right\}, \emptyset, \varphi^{\prime}, M\right)
\end{align*}
$$

After the goal has been reduced, the quantifiers must be eliminated from the formula. This elimination procedure is described in the next section.

## 5 Eliminating the Closed World

The formula $\exists M, \exists \varphi \beta_{G}\left(\varphi, M, \mathcal{I D}_{\mathcal{T}}\right)$ contains a second-order variable $M$ ranging over all the closed worlds, of which there are an infinite number. The next step in model generation is the elimination of the variable $M$. The elimination of $M$ means that we construct a new formula $\exists \varphi \beta_{G}^{\prime}(\varphi)$ that does not contain $M$, but the solutions to this formula can be used to construct a finite set of facts $M$ satisfying the original formula.

Elimination is accomplished by constructing a finite candidate set $M_{C}$ of non-ground terms with the property that there exists a satisfying $M$ if and only if there exists a subset $M_{C}^{\prime} \subseteq M_{C}$ where $\exists \varphi \beta_{G}\left(\varphi, \varphi\left(M_{C}^{\prime}\right), \mathcal{I} \mathcal{D}_{\mathcal{T}}\right)$. Once $M_{C}^{\prime}$ is discovered, all that remains is to arbitrarily choose an assignment of variables to constants to get a concrete $M$.

In the interest of space, we describe the solution when all backwards chaining paths pass through at most one negated derived term. The results here are easily generalized to arbitrary depth of negation. (The previous example fits into this restricted case.) Consider any expansion $\widehat{\beta}_{G}$ of the goal predicate. The results from the previous section guarantee that it has the following simplified form:

$$
\begin{align*}
& \exists M, \exists \varphi\left(\Theta^{\prime} \subseteq \operatorname{ker} \varphi\right) \wedge \alpha\left(P^{\prime}, N^{\prime}, \varphi, M\right) \wedge \forall \varphi^{\prime} \\
& {\left[\neg \omega_{1}\left(\varphi, \varphi^{\prime}, V_{1}^{+}\right) \vee\left(\Theta_{1}^{\prime} \nsubseteq \operatorname{ker} \varphi^{\prime}\right) \vee \neg \alpha\left(P_{1}^{\prime}, N_{1}^{\prime}, \varphi^{\prime}, M\right)\right] \wedge}  \tag{41}\\
& \left.\neg \omega_{2}\left(\varphi, \varphi^{\prime}, V_{2}^{+}\right) \vee\left(\Theta_{2}^{\prime} \nsubseteq \operatorname{ker} \varphi^{\prime}\right) \vee \neg \alpha\left(P_{2}^{\prime}, N_{2}^{\prime}, \varphi^{\prime}, M\right)\right] \wedge \\
& \vdots \\
& {\left[\neg \omega_{k}\left(\varphi, \varphi^{\prime}, V_{k}^{+}\right) \vee\left(\Theta_{k}^{\prime} \nsubseteq \operatorname{ker} \varphi^{\prime}\right) \vee \neg \alpha\left(P_{k}^{\prime}, N_{k}^{\prime}, \varphi^{\prime}, M\right)\right]}
\end{align*}
$$

The simplification results in exactly one non-negated acceptor $\alpha\left(P^{\prime}, N^{\prime}, \varphi, M\right)$, where $P^{\prime}$ and $N^{\prime}$ are the composition of many non-monotonic acceptors according to Lemma 3. Each negated derived term also creates a backwards chaining tree according to Equation 28 and these trees can be simplified in the same way. After simplification, the negative derived terms yield disjunctions of negated $\omega$ formulas, kernel constraints, and non-monotonic acceptors. These negated subformulas are arbitrarily numbered $1, \ldots, k$ and primed to remind the reader that they result from simplification.

Recall that every clause has a least one positive term in the tail, so every $P_{i}^{\prime}$ must be non-empty. However, the goal $G$ might not have positive terms, in which case $P^{\prime}=\emptyset$. If this holds, let $M=\emptyset$, then $M$ trivially satisfies $\alpha\left(\emptyset, N^{\prime}, \varphi, \emptyset\right)$ and trivial satisfies $\forall \varphi^{\prime}, \neg \alpha\left(P_{i}^{\prime}, N_{i}^{\prime}, \varphi^{\prime}, \emptyset\right)$. Thus, we focus on the interesting case where the goal $G$ contains some positive terms, i.e. $P^{\prime} \neq \emptyset$.

Assume $P^{\prime} \neq \emptyset$, then by definition of $\alpha$ it must be that $\varphi\left(P^{\prime}\right) \subseteq M$. The set $P^{\prime}$ must have a homomorphic image in any solution $M$. Therefore, let the candidate solution $M_{C}=P^{\prime}$. It may be that the negative part $N^{\prime}$ disallows some of these terms, but this can be discovered later. Next, consider the negated acceptors. It could be that some acceptor $\alpha\left(P_{i}^{\prime}, N_{i}^{\prime}, \varphi^{\prime}, M\right)$ is satisfied by the candidate model $M_{C}$. A necessary condition for this to occur is:

$$
\begin{equation*}
\exists \varphi^{\prime} \omega_{i}\left(\varphi, \varphi^{\prime}, V_{i}^{+}\right) \wedge\left(\Theta_{i}^{\prime} \subseteq \operatorname{ker} \varphi^{\prime}\right) \wedge\left(\varphi^{\prime}\left(P_{i}^{\prime}\right) \subseteq M_{C}\right) \tag{42}
\end{equation*}
$$

This situation is only problematic if no term from $N_{i}^{\prime}$ has a homomorphic image in $M_{C}$. These problematic situations are mitigated by expanding $M_{C}$ with the maximum number of negative terms: $M_{C}^{n e w}=M_{C}^{\text {old }} \cup \varphi^{\prime}\left(N_{i}^{\prime}\right)$. This expansion is performed for every possible $\varphi^{\prime}$, of which there are a finite number (if $M_{C}$ is already finite). If $N_{i}^{\prime}$ contains variables not found in $P_{i}^{\prime}$, then these variables are given new names.

Expanding $M_{C}$ may provide new opportunities for negated acceptors to fail (i.e. the acceptor evaluates to true, so its negation is false). These new opportunities must be identified and may require further expansion of $M_{C}$. The strategy is to add the maximum number of terms to $M_{C}$ that allow all the negated acceptors to succeed. For each negated acceptor assign an operator $C_{i}^{-}$:

$$
\begin{equation*}
C_{i}^{-}(\varphi, X)=\bigcup_{\varphi^{\prime}}\left\{\varphi^{\prime}\left(N_{i}^{\prime}\right) \mid \omega_{i}\left(\varphi, \varphi^{\prime}, V_{i}^{+}\right) \wedge\left(\Theta_{i}^{\prime} \subseteq \operatorname{ker} \varphi^{\prime}\right) \wedge\left(\varphi^{\prime}\left(P_{i}^{\prime}\right) \subseteq X\right)\right\} \cup X \tag{43}
\end{equation*}
$$

For the single non-negated acceptor $\alpha\left(P^{\prime}, N^{\prime}, \varphi, M\right)$ assign the operator $C^{+}$:

$$
\begin{equation*}
C^{+}(\varphi, X)=X \cup \varphi\left(P^{\prime}\right) \tag{44}
\end{equation*}
$$

The problem of finding the maximum $M_{C}$ can now be stated in terms of a least fixed-point equation:

$$
\begin{equation*}
M_{C}=C^{+}\left(\varphi, M_{C}\right)=C_{1}^{-}\left(\varphi, M_{C}\right)=C_{2}^{-}\left(\varphi, M_{C}\right)=\ldots=C_{k}^{-}\left(\varphi, M_{C}\right) \tag{45}
\end{equation*}
$$

These operators have two important properties: (1) monotonic: $X \subseteq C^{+/-}(X)$ (2) extensive: $X \subseteq Y \Rightarrow C^{+/-}(X) \subseteq C^{+/-}(Y)$. These properties lead to the important lemma:

Lemma 6. If the least fixed-point $M_{C}$ exists, then it is unique for a given $\varphi$.
In fact, $\varphi$ serves as a book-keeping mechanism to remember constraints over variables, and it can be constructed while solving the fixed-point equation. However, in the interest of space we omit the algorithm that constructs the fixed-point $M_{C}$.

Theorem 3. For some expansion $\widehat{\beta}_{G}$, if the fixed-point $M_{C}$ exists, then

$$
\begin{equation*}
\forall M, \forall \varphi \quad \widehat{\beta}_{G}(\varphi, M, \mathcal{I D}) \Rightarrow\left(\exists M_{C}^{\prime} \subseteq M_{C}, \exists \varphi^{\prime} \quad \varphi^{\prime}\left(M_{C}^{\prime}\right)=M\right) \tag{46}
\end{equation*}
$$

where $\varphi^{\prime}$ assigns variables to constants, and $M$ is a minimal solution.
This key theorem explains that if a least fixed-point $M_{C}$ exists for some expansion of the goal $\widehat{\beta}_{G}$, then a definite conclusion can be drawn about the satisfiability of this expansion: If there is a finite model that satisfies the subgoal, then there is some minimal model $M$ that satisfies the subgoal. The minimal model $M$ is exactly some subset $M_{C}^{\prime} \subseteq M_{C}$ where the variables of $M_{C}^{\prime}$ have been assigned to constants. The least fixed-point $M_{C}$ is finite, so there are a finite number of subsets $M_{C}^{\prime}$. Furthermore, there are only a finite number of interesting ways that variables can be assigned constants. This result leads to an algorithm for model generation:

```
Model Generation Algorithm
enumeration Results = { satisfiable, unsatisfiable, unknown };
let result := unsatisfiable;
Restratify(}(\mp@subsup{\Upsilon}{F}{},\mp@subsup{\Upsilon}{D}{},\Lambda,G)
let goal_expansions:= {\mp@subsup{\widehat{\beta}}{1}{},\mp@subsup{\widehat{\beta}}{2}{},\ldots,\mp@subsup{\widehat{\beta}}{k}{}}\mathrm{ ;}
for each }\mp@subsup{\widehat{\beta}}{i}{}\in\mathrm{ goal_expansions {
    if (fixed_point_exists( }\mp@subsup{\widehat{\beta}}{i}{})\mathrm{ ) {
        let M}\mp@subsup{M}{C}{}=\mathrm{ fixed_point ( }\mp@subsup{\widehat{\beta}}{i}{})\mathrm{ ;
        for each }\mp@subsup{M}{C}{\prime}\subseteq\mp@subsup{M}{C}{}\mathrm{ and each interesting }\mp@subsup{\varphi}{}{\prime}
            if ( \exists\varphi \widehat{\beta}
                result := satisfiable;
            return \varphi'( }\mp@subsup{M}{C}{\prime})\mathrm{ ;
        }
    }
    else result := unknown;
}
```

Fig. 3. Eliminating the closed world using fixed-points of goal expansions $\widehat{\beta}_{i}$.

### 5.1 Generating Schedules: Part 2

We apply these results to generate non-trivial models for the scheduling language. Recall that Equation 40 is the simplification of the goal predicate $\beta_{G}$ and this predicate has only one expansion to $\widehat{\beta}$. The task is the calculation of the fixed-point $M_{C}$ from $\widehat{\beta}$. Initially $M_{C}$ contains only the positive part of the goal:

$$
\begin{equation*}
M_{C}^{0}=\left\{\operatorname{processor}\left(p_{1}\right), \operatorname{processor}\left(p_{2}\right), \text { constraint }\left(t_{1}, t_{2}\right), \text { constraint }\left(t_{2}, t_{3}\right)\right\} \tag{47}
\end{equation*}
$$

This candidate model may violate the negated acceptors:

$$
\begin{aligned}
& \forall \varphi^{\prime} \neg \alpha\left(\left\{\operatorname{constraint}\left(x_{4}, y_{4}\right)\right\},\left\{\operatorname{task}\left(x_{4}\right)\right\}, \varphi^{\prime}, M\right) \wedge \\
& \forall \varphi^{\prime} \neg \alpha\left(\left\{\operatorname{constraint}\left(x_{5}, y_{5}\right)\right\},\left\{\operatorname{task}\left(y_{5}\right)\right\}, \varphi^{\prime}, M\right)
\end{aligned}
$$

The substitution $\varphi^{\prime}\left(x_{4}\right)=\varphi^{\prime}\left(x_{5}\right) \mapsto t_{1}, \varphi^{\prime}\left(y_{4}\right)=\varphi^{\prime}\left(y_{5}\right) \mapsto t_{2}$ is a witness to this possibility. (There exists a similar substitution for $t_{2}, t_{3}$.) These substitutions expand the candidate set to include the three tasks:

$$
\begin{equation*}
M_{C}^{1}=M_{C}^{0} \cup\left\{\operatorname{task}\left(t_{1}\right), \operatorname{task}\left(t_{2}\right), \operatorname{task}\left(t_{3}\right)\right\} \tag{48}
\end{equation*}
$$

This candidate set does not contain taskmap terms from tasks to processors, violating the subformula:

$$
\forall \varphi^{\prime} \neg \alpha\left(\left\{\operatorname{task}\left(x_{1}\right)\right\},\left\{\operatorname{taskmap}\left(x_{1}, y_{1}\right)\right\}, \varphi^{\prime}, M\right)
$$

The expansion of $M_{C}$ introduces the taskmap terms and also new variables:

$$
\begin{equation*}
M_{C}^{2}=M_{C}^{1} \cup\left\{\operatorname{taskmap}\left(t_{1}, x\right), \operatorname{taskmap}\left(t_{2}, y\right), \operatorname{taskmap}\left(t_{3}, z\right)\right\} \tag{49}
\end{equation*}
$$

Finally, new processor terms are introduced for the end-points of the taskmap terms:

$$
\begin{equation*}
M_{C}^{3}=M_{C}^{2} \cup\{\operatorname{processor}(x), \operatorname{processor}(y), \text { processor }(z)\} \tag{50}
\end{equation*}
$$

This set is the fixed-point, i.e. $M_{C}=M_{C}^{3}$.
The model generation problem can be solved by examining subsets of the fixed-point $M_{C}$. Some subsets will not produce satisfying models:

$$
M_{\text {fail }}=\left\{\begin{array}{c}
\operatorname{task}\left(t_{1}\right), \operatorname{task}\left(t_{2}\right), \operatorname{task}\left(t_{3}\right),  \tag{51}\\
\operatorname{taskmap}\left(t_{1}, x\right), \operatorname{taskmap}\left(t_{2}, y\right), \operatorname{taskmap}\left(t_{3}, z\right) \\
\operatorname{constraint}\left(t_{1}, t_{2}\right), \operatorname{constraint}\left(t_{2}, t_{3}\right), \operatorname{processor}\left(p_{1}\right),
\end{array}\right\}
$$

where $x=y=z=p_{1}$. This subset will fail for any assignment of variables to constants because there are not enough distinct processors. On the other hand, the set

$$
\begin{equation*}
M_{\text {success }}=M_{\text {fail }} \cup\left\{\operatorname{processor}\left(p_{2}\right)\right\} \tag{52}
\end{equation*}
$$

where $x=z=p_{1}$ and $y=p_{2}$ and $x \neq y$ satisfies the goal. In fact, any choice of constants that respects the equalities/disequalities satisfies the goal. In the next section we show how these subsets can be calculated using boolean satisfiability.

### 5.2 A Better Algorithm Using SMT

The simple algorithm in Figure 3 is a brute force approach for finding a model that satisfies the goal. It tries every subset of $M_{C}$ and every interesting assignment $\varphi^{\prime}$ of variables to constants. In general, there are an exponential number of $\left(M_{C}^{\prime}, \varphi^{\prime}\right)$ pairs to test. This exponential blow-up cannot be eliminated entirely, but it can be mitigated by translating the problem into a SAT problem. Mature SAT algorithms can be used to suppress the exponential blow-up. The encoding described here assumes a modern solver capable of reasoning about equalities among a set of non-boolean variables. We use an SMT solver (Satisfiability Modulo Theories) called Z3 [17] to accomplish this task. Z3 utilizes efficient SAT algorithms to solve problems that are not purely boolean; e.g. problems with equalities over non-boolean variables.

The first step of the encoding is the translation of complex terms to boolean variables. At this stage all the terms that need to be considered already exist in $M_{C}$, so the encoding is simple. Assign a boolean variable $\tau_{i}$ to each term $t_{i}$ in the candidate set $M_{C}$. If $\tau_{i}$ is true then the corresponding term $t_{i}$ is in the solution, otherwise $t_{i}$ is not in the solution. Furthermore, introduce a set of non-boolean variables $X$, which are the variables occurring as subterms in $M_{C}$ : $X=\operatorname{vars}\left(M_{C}\right)$. We provide these non-boolean variables so the SAT solver can decide if some variables $x_{i}, x_{j}$ in $M_{C}$ should take the the same values $\left(x_{i}=x_{j}\right)$, or different values $\left(x_{i} \neq x_{j}\right)$, or a fixed value $\left(x_{i}=\mathrm{c}\right)$.

The non-negated acceptor $\alpha\left(P^{\prime}, N^{\prime}, \varphi, M\right)$ defines terms $P^{\prime}$ that must be in any solution. Let $t_{1}^{+}, t_{2}^{+}, \ldots, t_{k}^{+}$be the terms in $M_{C}$ that were added according to $P^{\prime}$, then the following boolean formula must be true:

$$
\begin{equation*}
\bigwedge_{1 \leq j \leq k} \tau_{j}^{+} \tag{53}
\end{equation*}
$$

Next, consider any negated acceptor $\neg \alpha\left(P_{i}^{\prime}, N_{i}^{\prime}, \varphi^{\prime}, M\right)$. Let $\varphi^{\prime}$ be such that $\omega_{i}\left(\varphi, \varphi^{\prime}, V_{i}^{+}\right)$holds, $\Theta_{i}^{\prime} \subseteq \operatorname{ker} \varphi^{\prime}$, and $\varphi^{\prime}\left(P_{i}^{\prime}\right) \subseteq M_{C}$. Let $\varphi^{\prime}\left(P_{i}^{\prime}\right)=\left\{t_{1}^{-}, t_{2}^{-}, \ldots, t_{k}^{-}\right\}$. (Note $\varphi^{\prime}$ may not exist.) The negated acceptor is satisfied if one of the following holds:

1. One of the $t_{j}^{-}$terms is not in the solution.
2. There exists a pair of variables $(x, y) \in \operatorname{vars}\left(\varphi^{\prime}\left(P_{i}^{\prime}\right)\right)$ that is also in the kernel of $\varphi^{\prime}$, but $x \neq y$.
3. There exists a variable $x \in \operatorname{vars}\left(\varphi^{\prime}\left(P_{i}^{\prime}\right)\right)$ and a constant c where $(x, \mathrm{c}) \in$ ker $\varphi^{\prime}$, but $x \neq \mathrm{c}$.
4. There exists an extension of $\varphi^{\prime}$ to $\varphi^{\prime \prime}$ and some $t^{+} \in M_{C}$ so that $t^{+} \in \varphi^{\prime \prime}\left(N_{i}^{\prime}\right)$.

Conditions (1)-(3) yield the following encoding:

$$
\begin{equation*}
\left(\bigvee_{\substack{ \\1 \leq j \leq k}}^{\bigvee_{j}^{-}} \bigvee_{\substack{x, y \in \operatorname{vars}\left(\varphi^{\prime}\left(P_{i}^{\prime}\right)\right),(x, y) \in \operatorname{ker} \varphi^{\prime}}} x \neq y\right) \vee\left(\bigvee_{\substack{\operatorname{vars}\left(\varphi^{\prime}\left(P_{i}^{\prime}\right)\right), \mathrm{c} \in \Sigma,(x, \mathrm{c}) \in \operatorname{ker} \varphi^{\prime}}} x \neq \mathrm{c}\right) \tag{54}
\end{equation*}
$$

Condition (4) has a more complicated encoding:

$$
\begin{equation*}
\bigvee_{t^{+}}\left[\tau^{+} \wedge\left(\bigwedge_{\substack{\left.x, y \in \operatorname{vars}\left(\varphi^{\prime \prime}\left(P_{i}^{\prime}\right) \cup t^{+}\right)\right),(x, y) \in \operatorname{ker} \varphi^{\prime \prime}}} x=y\right) \wedge\left(\bigwedge_{\substack{\left.x \in \operatorname{vars}\left(\varphi^{\prime \prime}\left(P_{i}^{\prime}\right) \cup t^{+}\right)\right), c \in \Sigma,(x, \mathrm{c}) \in \operatorname{ker} \varphi^{\prime \prime}}} x=\mathrm{c}\right)\right] \tag{55}
\end{equation*}
$$

Such an encoding must be generated for all relevant $\varphi^{\prime}$ and $\varphi^{\prime \prime}$. A similar translation encodes the negative part $N^{\prime}$ of the non-negated acceptor $\alpha^{\prime}$. We omit it this in the interest of space.

In this discussion we have ignored the relationship between variables and terms. For example, two terms $f(x), f(y)$ are affected by equating (or disequating) the variables: If $x=y(x \neq y)$ then $f(x)=f(y)(f(x) \neq f(y))$. In this case the relationship can be easily encoded because all the variables in $M_{C}$ will only take constant values. Let $t_{i}, t_{j}$ be two terms that unify by equating variables to other variables or constants. Then the following must hold:

$$
\begin{equation*}
\left(\tau_{i} \Leftrightarrow \tau_{j}\right) \vee\left(\bigvee_{(x, y) \in m g u\left(t_{i}, t_{j}\right)} x \neq y\right) \vee\left(\bigvee_{(x, \mathrm{c}) \in m g u\left(t_{i}, t_{j}\right)} x \neq \mathrm{c}\right) \tag{56}
\end{equation*}
$$

Terms that unify by assigning some variables to complex terms are ignored.

## 6 Conclusion and Future Work

Model generation is an important tool for the model-based design of software systems. It can be used to generate non-trivial solution instances from domainspecific abstractions, perform design-space exploration, and reason about modeltransformations. We have given a sound algorithm that generates models from non-recursive and stratified Horn logic. These algorithms have been implemented in a tool called FORMULA (FORmal Modeling Using Logic Analysis). The SMT (SAT Modulo Theories) solver Z3 is used to solve the SAT encodings output by FORMULA.

Future work includes extending model generation to encompass constraint logic programming (CLP) frameworks. CLP combines Horn logic with constraints, as in the following clause:

$$
\text { bad_sched }(\text { critial_task }(x)) \leftarrow \operatorname{critical\_ task~}(x), \operatorname{task}(y), \operatorname{priority}(x)<\operatorname{priority}(y)
$$

Assume there is a new type of task called a critical_task. This clause states that a bad schedule assigns a high priority to a non-critical task. Priorities are expressed by the ordering < over integers, resulting in a combination of Horn logic with the theory of integers. We will utilize additional theories available at the SMT level to generate models for CLP extensions.

## 7 Acknowledgments

We would like to thank Nikolaj Bjørner for his invaluable feedback and his insight into Z3.

## References

1. Jackson, E.K., Sztipanovits, J.: Towards a formal foundation for domain specific modeling languages. Proceedings of the Sixth ACM International Conference on Embedded Software (EMSOFT'06) (October 2006) 53-62
2. Object Management Group: Mda guide version 1.0.1. Technical report (2003)
3. Bezivin, J., Gerbé, O.: Towards a precise definition of the omg/mda framework. In Proceedings of the 16th Conference on Automated Software Engineering (November 2001) 273-280
4. G. Karsai, J. Sztipanovits, A.L.T.B.: Model-integrated development of embedded software. Proceedings of the IEEE 91(1) (January 2003) 145-164
5. J. Burch, R. Passerone, A.S.V.: Modeling techniques in design-by-refinement methodologies. In: Integrated Design and Process Technology. (June 2002)
6. Lee, E.A., Neuendorffer, S.: Actor-oriented models for codesign: Balancing re-use and performance. Formal Methods and Models for Systems, Kluwer (2004)
7. Przymusinski, T.C.: Every logic program has a natural stratification and an iterated least fixed point model. In: PODS '89: Proceedings of the eighth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems, New York, NY, USA, ACM (1989) 11-21
8. Jensen, T.R., Toft, B.: Graph Coloring Problems. Wiley-Interscience, New York, ISBN 0-471-02865-7
9. Reiter, R.: On closed world data bases. In: Readings in nonmonotonic reasoning, Morgan Kaufmann Publishers Inc. (1987) 300-310
10. Blass, A., Gurevich, Y.: Existential fixed-point logic. In: Computation theory and logic, Springer-Verlag (1987) 20-36
11. van Gelder, A., Ross, K., Schlipf, J.S.: The well-founded semantics for general logic programs. Journal of the ACM 38(3) (1991) 620-650
12. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In Kowalski, R.A., Bowen, K., eds.: Proceedings of the Fifth International Conference on Logic Programming, Cambridge, Massachusetts, The MIT Press (1988) 10701080
13. Marek, V.W., Nerode, A., Remmel, J.B.: A context for belief revision: Forward chaining - normal nonmonotonic rule systems. Ann. Pure Appl. Logic 67(1-3) (1994) 269-323
14. Emden, M.H.V., Kowalski, R.A.: The semantics of predicate logic as a programming language. J. ACM 23(4) (1976) 733-742
15. Apt, K.R., Doets, K.: A new definition of SLDNF-resolution. The Journal of Logic Programming 18(2) (February 1994) 177-190
16. Dantsin, E., Eiter, T., Gottlob, G., Voronkov, A.: Complexity and expressive power of logic programming. ACM Comput. Surv. 33(3) (2001) 374-425
17. de Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In Proceedings of Fourteenth International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008) (March 2008)

[^0]:    ${ }^{1}$ We follow the convention that the OR of the empty set is false and the AND of the empty set is true.

