

On the generalization of learned constraints for ASP solving in temporal domains

JAVIER ROMERO

TORSTEN SCHAUB

KLAUS STRAUCH

University of Potsdam, Germany

submitted [n/a]; revised [n/a]; accepted [n/a]

Abstract

The representation of a dynamic problem in ASP usually boils down to using copies of variables and constraints, one for each time stamp, no matter whether it is directly encoded or via an action or temporal language. The multiplication of variables and constraints is commonly done during grounding and the solver is completely ignorant about the temporal relationship among the different instances. On the other hand, a key factor in the performance of today's ASP solvers is conflict-driven constraint learning. Our question is now whether a constraint learned for particular time steps can be generalized and reused at other time stamps, and ultimately whether this enhances the overall solver performance on temporal problems. Knowing full well the domain of time, we study conditions under which learned dynamic constraints can be generalized. We propose a simple translation of the original logic program such that, for the translated programs, the learned constraints can be generalized to other time points. Additionally, we identify a property of temporal problems that allows us to generalize all learned constraints to all time steps. It turns out that this property is satisfied by many planning problems. Finally, we empirically evaluate the impact of adding the generalized constraints to an ASP solver.

KEYWORDS: Answer Set Programming and Answer Set Solving and Temporal Reasoning

1 Introduction

Although Answer Set Programming (ASP; (Gelfond and Lifschitz 1988)) experiences an increasing popularity in academia and industry, a closer look reveals that this concerns mostly static domains. There is still quite a chasm between ASP's level of development for addressing static and dynamic domains. This is because its modeling language as well as its solving machinery aim so far primarily at static knowledge, while dynamic knowledge is mostly dealt with indirectly via reductions to the static case. This also applies to dedicated dynamic formalisms like action and temporal languages (Aguado et al. 2013; Gelfond and Lifschitz 1998). In fact, their reduction to ASP or SAT usually relies on translations that introduce a copy of each variable for each time step. The actual dynamics of the problem is thus compiled out and a solver treats the result as any other static problem.

We address this by proposing a way to (partly) break the opaqueness of the actual dynamic problem and equip an ASP solver with means for exploiting its temporal nature. More precisely, we introduce a method to strengthen the conflict-driven constraint learning framework (CDCL)

of ASP solvers so that dynamic constraints learned for specific time points can be generalized to other points in time. These additional constraints can in principle reduce the search space and improve the performance of the ASP solvers.

We start by reviewing some background material in Section 2. Next, in Section 3, we introduce a simple but general language to reason about time in ASP. We then define temporal problems, and characterize their solutions in terms of completion and loop nogoods, paralleling the approach to regular ASP solving (Gebser et al. 2012). In Section 4, using this language, we study conditions under which learned constraints can be generalized to other time steps. With it, in Section 5, we propose a simple translation such that, for the translated programs, learned constraints can be generalized to other time points without the need for any proof method. In addition, in Section 6 we identify a sufficient condition for the generalization of all learned constraints, that happens to hold for the planning problems that we encountered. Finally, in Section 7 we empirically evaluate the impact of adding the generalized constraints to the ASP solver *clingo*.

Our work can be seen as a continuation of the approach of *ginkgo* (Gebser et al. 2016), which also aimed at generalizing temporal constraints but resorted to an external inductive proof method (in ASP) for warranting correctness. More generally, a lot of work has been conducted over recent years on lazy ASP solving (Lefèvre et al. 2017; Palù et al. 2009; Weinzierl et al. 2020). Notably, conflict generalization was studied from a general perspective in (Comptoi-Taupe et al. 2020), dealing with several variables over heterogeneous domains. Lazy grounding via propagators was investigated in (Cuteri et al. 2020). Finally, it is worth mentioning that the usage of automata, as done in (Cabalar et al. 2021), completely abolishes the use of time points. A detailed formal and empirical comparative study of these approaches is interesting future work.

This is an extended version of the conference paper (Romero et al. 2022) presented at RuleML+RR 2022. The main new contribution is the identification in Section 6 of a property of temporal problems that allows us to generalize the learned constraints to all time points without the need of any translation. This improves significantly the applicability of our approach. In fact, this property is satisfied by the planning domains that we considered in our empirical evaluation of (Romero et al. 2022). Given this, we ran again those experiments, but this time with the original encodings, only slightly modified to satisfy the mentioned property. In addition to this, we have added the proofs to the theoretical results in Appendix A.

2 Background

We review the material from (Gebser et al. 2007) about solving normal logic programs, and adapt it for our purposes to cover normal logic programs with choice rules and integrity constraints over some set $\mathcal{P}$ of atoms.

A rule r has the form $H \leftarrow B$ where B is a set of literals over $\mathcal{P}$, and H is either an atom $p \in \mathcal{P}$, and we call r a *normal rule*, or $\{p\}$ for some atom $p \in \mathcal{P}$, making r a *choice rule*, or $\perp$, so that r is an *integrity constraint*. We usually drop braces from rule bodies B , and drop the arrow $\leftarrow$ when B is empty. We use the extended choice rule $\{p_1; \dots; p_n\} \leftarrow B$ as a shorthand for the choice rules $\{p_1\} \leftarrow B, \dots, \{p_n\} \leftarrow B$. A program Π is a set of rules. By Π^n , Π^c , and Π^i we denote its normal rules, choice rules and integrity constraints, respectively. Semantically, a logic program induces a collection of *stable models*, which are distinguished models of the program determined by the stable models semantics (see (Gebser et al. 2012; Gelfond and Lifschitz 1988) for details).

For a rule r of the form $H \leftarrow B$, let $h(r) = p$ be the *head* of r if H has the form p or $\{p\}$ for

some atom $p \in \mathcal{P}$, and let $h(r) = \perp$ otherwise. Let $B(r) = B$ be the *body* of r , $B(r)^+ = \{p \mid p \in \mathcal{P}, p \in B\}$ be the *positive body* of r , and $B(r)^- = \{p \mid p \in \mathcal{P}, \neg p \in B\}$ be the *negative body* of r . The set of atoms occurring in a rule r and in a logic program Π are denoted by $A(r)$ and $A(\Pi)$, respectively. The set of bodies in Π is $B(\Pi) = \{B(r) \mid r \in \Pi\}$. For regrouping rule bodies sharing the same head p , we define $B(p) = \{B(r) \mid r \in \Pi, h(r) = p\}$, and by $B^n(p)$ we denote the restriction of that set to bodies of normal rules, i.e., $\{B(r) \mid r \in \Pi^n, h(r) = p\}$.

A Boolean *assignment* A over a set $\mathcal{A}$, called the domain of A , is a set $\{\sigma_1, \dots, \sigma_n\}$ of *signed literals* σ_i of the form $\mathbf{T}p$ or $\mathbf{F}p$ for some $p \in \mathcal{A}$ and $1 \leq i \leq n$; $\mathbf{T}p$ expresses that p is *true* and $\mathbf{F}p$ that it is *false*. We omit the attribute *signed* for literals whenever clear from the context. We denote the complement of a literal σ by $\bar{\sigma}$, that is, $\overline{\mathbf{T}p} = \mathbf{F}p$ and $\overline{\mathbf{F}p} = \mathbf{T}p$. Given this, we access true and false propositions in A via $A^{\mathbf{T}} = \{p \in \mathcal{A} \mid \mathbf{T}p \in A\}$ and $A^{\mathbf{F}} = \{p \in \mathcal{A} \mid \mathbf{F}p \in A\}$. We say that a set of atoms X is consistent with an assignment A if $A^{\mathbf{T}} \subseteq X$ and $A^{\mathbf{F}} \cap X = \emptyset$. In our setting, a *nogood* is a set $\{\sigma_1, \dots, \sigma_n\}$ of signed literals, expressing a constraint violated by any assignment containing $\sigma_1, \dots, \sigma_n$. Accordingly, the nogood for a body B , denoted by $ng(B)$, is $\{\mathbf{T}p \mid p \in B^+\} \cup \{\mathbf{F}p \mid p \in B^-\}$. We say that an assignment A over $\mathcal{A}$ is *total* if $A^{\mathbf{T}} \cup A^{\mathbf{F}} = \mathcal{A}$ and $A^{\mathbf{T}} \cap A^{\mathbf{F}} = \emptyset$. A total assignment A over $\mathcal{A}$ is a *solution* for a set Δ of nogoods, if $\delta \not\subseteq A$ for all $\delta \in \Delta$. A set Δ of nogoods *entails* a nogood δ if $\delta \not\subseteq A$ for all solutions A over $\mathcal{A}$ for Δ , and it entails a set of nogoods ∇ if it entails every nogood $\delta \in \nabla$ in the set.

We say that a nogood δ is a *resolvent* of a set of nogoods Δ if there is a sequence of nogoods $\delta_1, \dots, \delta_n$ with $n \geq 1$ such that $\delta_n = \delta$, and for all i such that $1 \leq i \leq n$, either $\delta_i \in \Delta$, or there are some δ_j, δ_k with $1 \leq j < k < i$ such that $\delta_i = (\delta_j \setminus \{\sigma\}) \cup (\delta_k \setminus \{\bar{\sigma}\})$ for some signed literal σ . In this case, we say that the sequence $\delta_1, \dots, \delta_n$ is a *proof* of δ_n . We say that a signed literal σ is *unit resulting* for a nogood δ and an assignment A if $\delta \setminus A = \{\sigma\}$ and $\bar{\sigma} \notin A$. For a set of nogoods Δ and an assignment A , unit propagation is the process of extending A with unit-resulting literals until no further literal is unit resulting for any nogood in Δ .

Inferences in ASP can be expressed in terms of atoms and rule bodies. We begin with nogoods capturing inferences from the Clark completion. For a body $\beta = \{p_1, \dots, p_m, \neg p_{m+1}, \dots, \neg p_n\}$, we have that $\delta(\beta) = \{\mathbf{F}\beta, \mathbf{T}p_1, \dots, \mathbf{T}p_m, \mathbf{F}p_{m+1}, \dots, \mathbf{F}p_n\}$ and $\Delta(\beta) = \{\{\mathbf{T}\beta, \mathbf{F}p_1\}, \dots, \{\mathbf{T}\beta, \mathbf{F}p_m\}, \{\mathbf{T}\beta, \mathbf{T}p_{m+1}\}, \dots, \{\mathbf{T}\beta, \mathbf{T}p_n\}\}$. For an atom p such that $B^n(p) = \{\beta_1, \dots, \beta_k\}$, we have that $\Delta(p) = \{\{\mathbf{F}p, \mathbf{T}\beta_1\}, \dots, \{\mathbf{F}p, \mathbf{T}\beta_k\}\}$, and if $B(p) = \{\beta_1, \dots, \beta_k\}$ then $\delta(p) = \{\mathbf{T}p, \mathbf{F}\beta_1, \dots, \mathbf{F}\beta_k\}$. Given this, the *completion nogoods* of a logic program Π are defined as follows:

$$\begin{aligned} \Delta_{\Pi} = & \{\delta(\beta) \mid \beta \in B(\Pi \setminus \Pi^i)\} \cup \{\delta \in \Delta(\beta) \mid \beta \in B(\Pi \setminus \Pi^i)\} \\ & \cup \{\delta(p) \mid p \in A(\Pi)\} \cup \{\delta \in \Delta(p) \mid p \in A(\Pi)\} \\ & \cup \{ng(\beta) \mid \beta \in B(\Pi^i)\} \end{aligned}$$

Choice rules of the form $\{p\} \leftarrow \beta$ are considered by not adding the corresponding nogood $\{\mathbf{F}p, \mathbf{T}\beta\}$ to $\Delta(p)$, and integrity constraints from Π^i of the form $\perp \leftarrow \beta$ are considered by adding directly their corresponding nogood $ng(\beta)$. The definition of the *loop nogoods* Λ_{Π} , capturing the inferences from loop formulas, is the same as in (Gebser et al. 2007). We do not specify them here since they do not pose any special challenge to our approach, and they are not needed in our (tight) examples.

To simplify the presentation, we slightly deviate from (Gebser et al. 2007) and consider a version of the nogoods of a logic program where the occurrences of the empty body are simplified. Note that $\delta(\emptyset) = \{\mathbf{F}\emptyset\}$ and $\Delta(\emptyset) = \emptyset$. Hence, if $\emptyset \in B(\Pi)$ then any solution to the completion and loop nogoods of Π must contain $\mathbf{T}\emptyset$. Based on this, we can delete from $\Delta_{\Pi} \cup \Lambda_{\Pi}$ the nogoods

that contain $\mathbf{F}\emptyset$, and eliminate the occurrences of $\mathbf{T}\emptyset$ from the others. Formally, we define the set of (simplified) nogoods for Π as:

$$\Sigma_{\Pi} = \{\delta \setminus \{\mathbf{T}\emptyset\} \mid \delta \in \Delta_{\Pi} \cup \Lambda_{\Pi}, \mathbf{F}\emptyset \notin \delta\}.$$

To accommodate this change, for a program Π , we fix the domain $\mathcal{A}$ of the assignments to the set $A(\Pi) \cup (B(\Pi) \setminus \emptyset)$. Given this, the stable models of a logic program Π can be characterized by the nogoods Σ_{Π} for that program. This is made precise by the following theorem, which is an adaptation of Theorem 3.4 from (Gebser et al. 2007) to our setting.

Theorem 1

Let Π be a logic program. Then, $X \subseteq A(\Pi)$ is a stable model of Π iff $X = A^{\mathbf{T}} \cap A(\Pi)$ for a (unique) solution A for Σ_{Π} .

To compute the stable models of a logic program Π , we apply the algorithm *CDNL-ASP*(Π) from (Gebser et al. 2007) implemented in the ASP solver *clingo*. The algorithm searches for a solution A to the set of nogoods Σ_{Π} , and when it finds one it returns the corresponding set of atoms $A^{\mathbf{T}} \cap A(\Pi)$. *CDNL-ASP* maintains a current assignment A and a current set of learned nogoods ∇ , both initially empty. The main loop of the algorithm starts by applying unit propagation to $\Sigma_{\Pi} \cup \nabla$, possibly extending A . Every derived literal is “implied” by some nogood $\delta \in \Sigma_{\Pi} \cup \nabla$, which is stored in association with the derived literal. This derivation may lead to the violation of another nogood. This situation is called *conflict*. If propagation finishes without conflict, then a (heuristically chosen) literal can be added to A , provided that A is partial, while otherwise A represents a solution and can be directly returned. On the other hand, if there is a conflict, there are two possibilities. Either it is a top-level conflict, independent of heuristically chosen literals, in which case the algorithm returns *unsatisfiable*. Or, if that is not the case, the conflict is analyzed to calculate a conflict nogood δ , that is added to ∇ . More in detail, δ is a resolvent of the set of nogoods associated with the literals derived after the last heuristic choice. Hence, every learned nogood δ added to ∇ is a resolvent of $\Sigma_{\Pi} \cup \nabla$ and, by induction, it is also a resolvent of Σ_{Π} . After recording δ , the algorithm backjumps to the earliest stage where the complement of some formerly assigned literal is implied by δ , thus triggering propagation and starting the loop again.

This algorithm has been extended for solving under assumptions (Eén and Sörensson 2003). In this setting, the procedure *CDNL-ASP*(Π, S) receives additionally as input a partial assignment S over $A(\Pi)$, the so-called assumptions, and returns some stable model of Π that is consistent with S . To accommodate this extension, the algorithm simply decides first on the literals from S , and returns *unsatisfiable* as soon as any of these literals is undone by backjumping. No more changes are needed. Notably, the learned nogoods are still resolvents of Δ_{Π} , that are independent of the set of assumptions S .

3 Temporal programs, problems and nogoods

We introduce a simple language of temporal logic programs to represent temporal problems. These programs represent the dynamics of a temporal domain by referring to two time steps: the current step and the previous step. We refer to the former by atoms from a given set $\mathcal{P}$, and to the latter by atoms from the set $\mathcal{P}' = \{p' \mid p \in \mathcal{P}\}$, that we assume to be disjoint from $\mathcal{P}$. Following the common-sense flow of time, normal or choice rules define the atoms of the current step in terms of the atoms of both the current and the previous step. Integrity constraints forbid some current states, possibly depending on the previous state. Syntactically, a temporal logic program Π over

$\mathcal{P}$ has the form of a (non-temporal) logic program over $\mathcal{P} \cup \mathcal{P}'$ such that for every rule $r \in \Pi$, if $r \in \Pi^n \cup \Pi^c$ then $h(r) \in \mathcal{P}$, and otherwise $(B(r)^+ \cup B(r)^-) \cap \mathcal{P} \neq \emptyset$. Given that temporal logic programs over $\mathcal{P}$ can also be seen as (non-temporal) logic programs over $\mathcal{P} \cup \mathcal{P}'$, in what follows we may apply the notation of the latter to the former. We say that the rules $r \in \Pi$ such that $A(r) \subseteq \mathcal{P}$ are static, and otherwise we say that they are dynamic.

One of the goals of the design of this language was to capture the core of the translations to ASP of action and temporal languages (Gelfond and Lifschitz 1998; Aguado et al. 2013). We do not elaborate this further, but from this perspective, temporal programs can be seen as an intermediate language in the workflow of ASP solving for those higher level languages. On the other hand, a variant of this language was used recently to represent the transition function of various types of planning problems (Fandinno et al. 2021). More in detail, this representation consists of choice rules of the form $\{a\} \leftarrow$ to generate the occurrences of actions a , normal rules of the form $f \leftarrow B$ to define the value of the fluents f in terms of the values of other fluents or actions at the current or previous steps, and integrity constraints of the form $\perp \leftarrow B$, where some action a belongs to B , to specify the preconditions of the actions.

Example 1

Our running example is the temporal logic program Π_1 over $\mathcal{P}_1 = \{a, b, c, d\}$ that consists only of choice rules and integrity constraints:

$$\begin{array}{ll} \{a; b; c; d\} \leftarrow & \perp \leftarrow a', \neg b \\ \perp \leftarrow \neg b', b & \perp \leftarrow \neg c', a \\ \perp \leftarrow d', b & \perp \leftarrow c, \neg d \\ \perp \leftarrow \neg a', \neg c & \perp \leftarrow \neg a', c', \neg a \end{array}$$

Temporal logic programs Π can be instantiated to specific time intervals. We introduce some notation for that. Let m and n be integers such that $1 \leq m \leq n$, and $[m, n]$ denote the set of integers $\{i \mid m \leq i \leq n\}$. For $p \in \mathcal{P}$, the symbol $p[m]$ denotes the atom p_m , and for $p' \in \mathcal{P}'$, the symbol $p'[m]$ denotes the atom p_{m-1} . For a set of atoms $X \subseteq \mathcal{P} \cup \mathcal{P}'$, $X[m]$ denotes the set of atoms $\{p[m] \mid p \in X\}$, and $X[m, n]$ denotes the set of atoms $\{p[i] \mid p \in X, i \in [m, n]\}$. For a rule r over $\mathcal{P} \cup \mathcal{P}'$, the symbol $r[m]$ denotes the rule that results from replacing in r every atom $p \in \mathcal{P} \cup \mathcal{P}'$ by $p[m]$, and $r[m, n]$ denotes the set of rules $\{r[i] \mid i \in [m, n]\}$. Finally, for a temporal program Π , $\Pi[m]$ is $\{r[m] \mid r \in \Pi\}$, and $\Pi[m, n]$ is $\{r[i] \mid r \in \Pi, i \in [m, n]\}$.

Example 2

The instantiation of Π_1 at 1, denoted by $\Pi_1[1]$, is:

$$\begin{array}{ll} \{a_1; b_1; c_1; d_1\} \leftarrow & \perp \leftarrow a_0, \neg b_1 \\ \perp \leftarrow \neg b_0, b_1 & \perp \leftarrow \neg c_0, a_1 \\ \perp \leftarrow d_0, b_1 & \perp \leftarrow c_1, \neg d_1 \\ \perp \leftarrow \neg a_0, \neg c_1 & \perp \leftarrow \neg a_0, c_0, \neg a_1 \end{array}$$

The programs $\Pi_1[i]$ for $i \in \{2, 3, 4\}$ are the same, except that the subindex 1 is replaced by i , and the subindex 0 is replaced by $i - 1$. The instantiation of Π_1 at $[1, 4]$, denoted by $\Pi_1[1, 4]$, is $\Pi_1[1] \cup \Pi_1[2] \cup \Pi_1[3] \cup \Pi_1[4]$.

To represent temporal reasoning problems, temporal programs are complemented by assignments I and F that partially or completely describe the initial and the final state of a problem. Formally, a *temporal logic problem* over some set of atoms $\mathcal{P}$ is a tuple $\langle \Pi, I, F \rangle$ where Π is a temporal logic program over $\mathcal{P}$, and I and F are assignments over $\mathcal{P}$. A solution to such a

problem is a sequence of states that is consistent with the dynamics described by Π and with the information provided by I and F . The possible sequences of states of length n , for some integer $n \geq 1$, are represented by the *generator program* for Π and n , denoted by $gen(\Pi, n)$, that consists of the rules $\{\{p_0\} \leftarrow \mid p \in \mathcal{P}\} \cup \Pi[1, n]$. Then, a *solution* to a temporal problem $\langle \Pi, I, F \rangle$ is defined as a pair (X, n) , where n is an integer such that $n \geq 1$, and X is a stable model of $gen(\Pi, n)$ consistent with $I[0] \cup F[n]$.

Temporal problems can be used to formalize planning problems, using a temporal logic program Π of the form described above, a total assignment I that assigns a value to every possible atom (action occurrences are made false initially), and a partial assignment F to fix the goal. The solutions of the temporal problem correspond to the plans of the planning problem.

Example 3

The temporal problem $\langle \Pi_1, \emptyset, \emptyset \rangle$ has three solutions of length 4: $(Y, 4)$, $(Y \cup \{d_2\}, 4)$, and $(Y \cup \{b_3\}, 4)$, where Y is the set of atoms $\{a_0, b_0, c_0, a_1, b_1, b_2, c_3, d_3, a_4, c_4, d_4\}$.

To pave the way to the nogood characterization of temporal logic problems, we define the transition program $trans(\Pi)$ of a temporal logic program Π as the (non-temporal) logic program $\Pi \cup \{\{p'\} \leftarrow \mid p' \in \mathcal{P}'\}$ over $\mathcal{P} \cup \mathcal{P}'$. Each stable model of this program represents a possible transition between a previous and a current step, where the former is selected by the additional choice rules over atoms from $\mathcal{P}'$, and the latter is determined by the rules of Π , interpreted as non-temporal rules.

Example 4

The transition program $trans(\Pi_1)$ is the (non-temporal) program $\Pi_1 \cup \{\{a'; b'; c'; d'\} \leftarrow\}$ over $\mathcal{P}_1 \cup \{p' \mid p \in \mathcal{P}_1\}$. Some stable models of $trans(\Pi_1)$ are $\{a', b', c', a, b\}$ and $\{c', d', a, c, d\}$, that correspond to the transitions to step 1 and step 4 of the solution $(Y, 4)$, respectively.

Next, we introduce temporal nogoods and their instantiation. Given a temporal logic program Π over $\mathcal{P}$, a temporal nogood over $\mathcal{P} \cup B(\Pi)$ has the form of a (non-temporal) nogood over $\mathcal{P} \cup \mathcal{P}' \cup B(\Pi)$. For a temporal nogood δ over $\mathcal{P} \cup B(\Pi)$ and an integer $n \geq 1$, the instantiation of δ at n , denoted by $\delta[n]$, is the nogood that results from replacing in δ any signed literal $\mathbf{T}\alpha$ ($\mathbf{F}\alpha$) by $\mathbf{T}\alpha[n]$ (by $\mathbf{F}\alpha[n]$, respectively). We extend this notation to sets of nogoods and to intervals like we did above. For example, $\delta_1 = \{\mathbf{F}b', \mathbf{T}b\}$ is a temporal nogood over $\mathcal{P}_1 \cup B(\Pi_1)$, and $\delta_1[1, 2]$ is $\{\{\mathbf{F}b_0, \mathbf{T}b_1\}, \{\mathbf{F}b_1, \mathbf{T}b_2\}\}$. By $step(\delta)$ we denote the interval of the steps of the literals occurring in δ , i.e., $step(\delta) = [\min(index(\delta)), \max(index(\delta))]$ where $index(\delta) = \{i \mid \mathbf{T}p_i \in \delta \text{ or } \mathbf{F}p_i \in \delta\}$. For example, $step(\{\mathbf{F}b_0, \mathbf{T}b_2\}) = [0, 2]$.

We are now ready to define the temporal nogoods for a temporal logic program Π over $\mathcal{P}$. Recall that $trans(\Pi)$ is a (non-temporal) logic program over $\mathcal{P} \cup \mathcal{P}'$, whose corresponding nogoods are denoted by $\Sigma_{trans(\Pi)}$. Then, the set of *temporal nogoods* for Π , denoted by Ψ_Π , has the form $\Sigma_{trans(\Pi)}$, interpreted as a set of temporal nogoods over $\mathcal{P} \cup B(\Pi)$, and not as a set of (non-temporal) nogoods over $\mathcal{P} \cup \mathcal{P}' \cup B(\Pi)$.

Example 5

The set Ψ_{Π_1} of temporal nogoods for Π_1 is $\{\{\mathbf{T}a', \mathbf{F}b\}, \{\mathbf{F}b', \mathbf{T}b\}, \{\mathbf{F}c', \mathbf{T}a\}, \{\mathbf{T}d', \mathbf{T}b\}, \{\mathbf{T}c, \mathbf{F}d\}, \{\mathbf{F}a', \mathbf{F}c\}, \{\mathbf{F}a', \mathbf{T}c', \mathbf{F}a\}\}$.

Temporal nogoods provide an alternative characterization of the nogoods of $gen(\Pi, n)$.

Proposition 1

Let Π be a temporal logic program, and $n \geq 1$ some integer. Then, $\Sigma_{gen(\Pi, n)} = \Psi_\Pi[1, n]$.

In words, the nogoods for $gen(\Pi, n)$ are the same as the instantiation of the temporal nogoods for Π , that are nothing else than the nogoods of the logic program $trans(\Pi)$ interpreted as temporal nogoods. Then, by Theorem 1, the temporal nogoods can be used to characterize the solutions of temporal logic problems.

Theorem 2

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem. The pair (X, n) is a solution to $\langle \Pi, I, F \rangle$ for some integer $n \geq 1$ and $X \subseteq \mathcal{P}[0, n]$ iff $X = A^T \cap \mathcal{P}[0, n]$ for a (unique) solution A for $\Psi_\Pi[1, n]$ such that $I[0] \cup F[n] \subseteq A$.

4 Generalization of learned constraints

A common software architecture to solve a temporal problem $\langle \Pi, I, F \rangle$ combines a scheduler that assigns resources to different values of n , with one or many solvers that look for solutions of the assigned lengths n (see (Rintanen et al. 2006), for example). The standard approach for the solvers is to extend the program $gen(\Pi, n)$ with facts and integrity constraints to adequately represent I and F , and call the procedure *CDNL-ASP* with this extended program without assumptions. This method does not work well for our purposes, because it leads to a nogood representation of the initial and the final steps that is different from the nogood representation of the other steps. Hence, the constraints learned using nogoods specific to the initial and final steps may not be generalizable to the other steps. To overcome this issue, in our approach the solvers apply the procedure $CDNL-ASP(gen(\Pi, n), I[0] \cup F[n])$ to the generator program for Π and n , using assumptions to fix the assignments about the initial and final situations. Observe that in this case, by Proposition 1, the solver initially contains exactly the nogoods $\Psi_\Pi[1, n]$, and all the nogoods that it learns afterwards are resolvents of $\Psi_\Pi[1, n]$.

Once this is settled, we ask ourselves:

What generalizations of the nogoods learned by CDNL-ASP can be applied to the same or other problems?

We make the question more precise step by step. First, instead of talking about “the nogoods learned by the algorithm”, we refer to the resolvents of $\Psi_\Pi[1, n]$ for some temporal problem $\langle \Pi, I, F \rangle$. Or more precisely, we refer to the resolvents of $\Psi_\Pi[i, j]$ for some i and j such that $1 \leq i \leq j \leq n$, since the learned nogoods are always the result of resolving nogoods belonging to some interval $[i, j]$ that may be smaller than $[1, n]$.

To formalize the notion of the “generalizations of nogoods”, we introduce some notation for shifting a non-temporal nogood an amount of t time steps. For integers $n \geq 1$ and t , and a non-temporal nogood δ over $(\mathcal{P} \cup \mathcal{P}' \cup B(\Pi))[1, n]$, the symbol $\delta\langle t \rangle$ denotes the nogood that results from replacing in δ any signed literal $\mathbf{T}\alpha_m$ ($\mathbf{F}\alpha_m$) by $\mathbf{T}\alpha_{m+t}$ (by $\mathbf{F}\alpha_{m+t}$, respectively). For example, $\delta\langle 0 \rangle = \delta$, and if $\delta = \{\mathbf{T}a_2, \mathbf{F}b_3\}$, then $\delta\langle 1 \rangle$ is $\{\mathbf{T}a_3, \mathbf{F}b_4\}$, and $\delta\langle -1 \rangle$ is $\{\mathbf{T}a_1, \mathbf{F}b_2\}$. We say that $\delta\langle t \rangle$ is a *shifted version* of the nogood δ , and that a *generalization* of a nogood is a set of some of its shifted versions. For example, $\{\{\mathbf{T}a_2, \mathbf{F}b_3\}\}$ and $\{\{\mathbf{T}a_1, \mathbf{F}b_2\}, \{\mathbf{T}a_2, \mathbf{F}b_3\}, \{\mathbf{T}a_3, \mathbf{F}b_4\}\}$ are generalizations of $\{\mathbf{T}a_2, \mathbf{F}b_3\}$ and of $\{\mathbf{T}a_3, \mathbf{F}b_4\}$.

Next, by the “other problems” mentioned in the question, we refer to variations m of the length of the solution, and to variations $\langle \Pi, I', F' \rangle$ of the original problem where the initial and final situation may change, but the temporal program remains the same. Then, a generalization of a nogood




Fig. 1. Representation of different shifted versions of the nogood $\delta = \{\mathbf{T}a\}[3]$. The surrounding rectangles cover the interval of the nogoods needed to prove them. For example, the rectangle of $\{\mathbf{T}a\}[2]$ covers the interval $[1, 3]$ because $\{\mathbf{T}a\}[2]$ is a resolvent of $\Psi_{\Pi_1}[1, 3]$.

“can be applied” to such problems if it can be added to the set of nogoods used by the algorithm *CDNL-ASP* without changing the solutions to the problem. For any variation $\langle \Pi, I', F' \rangle$, those nogoods are $\Psi_{\Pi}[1, m]$, and a generalization can be added to them if the generalization is entailed by them. Hence, a generalization of a nogood “can be applied” to “some problem” $\langle \Pi, I', F' \rangle$, searching for a solution of length m , if the generalization is entailed by $\Psi_{\Pi}[1, m]$. Putting all together, we can rephrase our question as follows:

Given some temporal logic problem $\langle \Pi, I, F \rangle$, what generalizations of a resolvent δ of $\Psi_{\Pi}[i, j]$ are entailed by $\Psi_{\Pi}[1, m]$?

Example 6

Consider a call of *CDNL-ASP*($gen(\Pi_1, n), \emptyset$) to search for a solution of length n to the temporal problem $\langle \Pi_1, \emptyset, \emptyset \rangle$, where n has the value 4. Initially, the solver may choose to make a_3 true by adding $\mathbf{T}a_3$ to the initial assignment. Then, by unit propagation, it could derive the literal $\mathbf{T}c_2$ by $\{\mathbf{F}c', \mathbf{T}a\}[3]$, the literal $\mathbf{T}d_2$ by $\{\mathbf{T}c, \mathbf{F}d\}[2]$, the literal $\mathbf{F}b_3$ by $\{\mathbf{T}d', \mathbf{T}b\}[3]$, and the literal $\mathbf{F}b_4$ by $\{\mathbf{F}b', \mathbf{T}b\}[4]$, leading to a conflict due to the violation of the nogood $\{\mathbf{T}a', \mathbf{F}b\}[4]$. At this stage, the solver would learn the nogood $\delta = \{\mathbf{T}a\}[3]$ by resolving iteratively $\{\mathbf{T}a', \mathbf{F}b\}[4]$ with the nogoods $\{\mathbf{F}b', \mathbf{T}b\}[4]$, $\{\mathbf{T}d', \mathbf{T}b\}[3]$, $\{\mathbf{T}c, \mathbf{F}d\}[2]$, and $\{\mathbf{F}c', \mathbf{T}a\}[3]$ used for propagation. Hence, δ is a resolvent of the set of those nogoods. Moreover, given that those nogoods are instantiations of some temporal nogoods of Ψ_{Π_1} at the interval $[2, 4]$, δ is also a resolvent of $\Psi_{\Pi_1}[2, 4]$ and of $\Psi_{\Pi_1}[1, n]$. Observe that, by shifting the nogoods -1 time points, we obtain that $\delta\langle -1 \rangle = \{\mathbf{T}a\}[2]$ is a resolvent of $\Psi_{\Pi_1}[1, 3]$, and therefore also of $\Psi_{\Pi_1}[1, n]$. Then, by the correctness of resolution, we have that the generalization $\{\{\mathbf{T}a\}[2], \{\mathbf{T}a\}[3]\}$ of δ is entailed by $\Psi_{\Pi_1}[1, n]$. On the other hand, $\delta\langle -2 \rangle = \{\mathbf{T}a\}[1]$ is a resolvent of $\Psi_{\Pi_1}[0, 2]$, but not of $\Psi_{\Pi_1}[1, n]$, (partly) because the instantiations at 0 do not belong to $\Psi_{\Pi_1}[1, n]$. Similarly, $\delta\langle 1 \rangle = \{\mathbf{T}a\}[4]$ is a resolvent of $\Psi_{\Pi_1}[3, 5]$, but not of $\Psi_{\Pi_1}[1, n]$, (partly) because the instantiations at 5 do not belong to $\Psi_{\Pi_1}[1, n]$ (see Figure 1).

This example suggests a sufficient condition for the generalization of a nogood δ learned from $\Psi_{\Pi}[i, j]$: a shifted version $\delta\langle t \rangle$ of some generalization of δ is entailed by $\Psi_{\Pi}[1, n]$ if the nogoods that result from shifting $\Psi_{\Pi}[i, j]$ an amount of t time points belong to $\Psi_{\Pi}[1, n]$. We answer our previous question by stating this condition precisely in the next theorem.

Theorem 3

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem, and δ be a resolvent of $\Psi_{\Pi}[i, j]$ for some i and j such

that $1 \leq i \leq j$. Then, for every $n \geq 1$, the set of nogoods $\Psi_\Pi[1, n]$ entails the generalization

$$\{\delta\langle t \rangle \mid [i + t, j + t] \subseteq [1, n]\}.$$

The proof is based on the fact that the resolution proof that derived δ from $\Psi_\Pi[i, j]$ can be used to derive every $\delta\langle t \rangle$ from $\Psi_\Pi[i + t, j + t]$, simply by shifting the nogoods t time steps. This means that $\delta\langle t \rangle$ is a resolvent of $\Psi_\Pi[i + t, j + t]$. Given that $[i + t, j + t] \subseteq [1, n]$, the nogood $\delta\langle t \rangle$ is also a resolvent of $\Psi_\Pi[1, n]$. Then, the theorem follows from the correctness of resolution.

This result allows us to generalize the learned nogoods to different lengths and different initial and final situations, as long as the specified conditions hold. Following our example, if we were now searching for a solution of length 9 to the temporal problem $\langle \Pi_1, \{\mathbf{T}c\}, \{\mathbf{T}b\} \rangle$, we could add the generalization $\{\{\mathbf{T}a\}[i] \mid i \in [2, 8]\}$ to $CDNL\text{-}ASP(gen(\Pi_1, 9), \{\mathbf{T}c_0, \mathbf{T}b_9\})$.

The theorem can be applied in an online setting, where the generalizations are added while solving, or in an offline setting, where the generalizations are stored to apply them later to other problems. Observe that to benefit the most from the result, we should know what is the specific interval $[i, j]$ of the nogoods used to obtain a learned nogood. We could obtain this information by modifying the solving algorithm, and recording that interval for every learned nogood. We leave that option for future work, and in the next section we follow another approach that does not require to modify the solver.

5 Program translations

In this section, we present a translation of the original temporal program such that the nogoods learned using the translated program can be generalized to all time points. We start with a simple translation tr^λ that works for temporal programs where all dynamic rules are integrity constraints. Later, we show that all temporal programs can be translated to this form.

We say that a temporal logic program Π over $\mathcal{P}$ is in *previous normal form* (PNF) if $A(\Pi \setminus \Pi^i) \cap \mathcal{P}' = \emptyset$, and that a temporal logic problem $\langle \Pi, I, F \rangle$ over $\mathcal{P}$ is in PNF if Π is in PNF. Given a temporal logic program Π over $\mathcal{P}$, let Π^{di} denote the set $\{r \mid r \in \Pi^i, A(r) \cap \mathcal{P}' \neq \emptyset\}$ of dynamic integrity constraints of Π . Note that if Π is in PNF, then the dynamic rules of Π belong to Π^{di} . The translation $tr^\lambda(\Pi)$ tags the rules in Π^{di} with a new atom λ , that does not belong to $\mathcal{P}$ or $\mathcal{P}'$, and extends the program with a choice rule for λ . Formally, by $tr^\lambda(\Pi)$ we denote the temporal logic program:

$$\Pi \setminus \Pi^{di} \cup \{\{\lambda\} \leftarrow\} \cup \{\perp \leftarrow B(r) \cup \{\lambda\} \mid r \in \Pi^{di}\}.$$

It is easy to see that when λ is chosen to be true, $tr^\lambda(\Pi)$ generates the same transitions as Π . Then, we can solve temporal programs $\langle \Pi, I, F \rangle$ by solving temporal problems $\langle tr^\lambda(\Pi), I, F \rangle$, if we consider only solutions that make λ true at all steps after the initial one. For convenience, at the initial step we consider only the case where λ is false. To make this precise, we say that a solution (X, n) to a temporal problem is λ -normal if $X \cap (\{\lambda\}[0, n]) = \{\lambda\}[1, n]$. The next proposition states the relation between these λ -normal solutions and the original solutions using Π .

Proposition 2

Let $\mathcal{T}_1 = \langle \Pi, I, F \rangle$ and let $\mathcal{T}_2 = \langle tr^\lambda(\Pi), I, F \rangle$ be temporal logic problems. There is a one-to-one correspondence between the solutions to $\mathcal{T}_1$ and the λ -normal solutions to $\mathcal{T}_2$.

The call $CDNL\text{-}ASP(gen(tr^\lambda(\Pi), n), I[0] \cup F[n] \cup \{\mathbf{F}\lambda_0\} \cup \{\mathbf{T}\lambda\}[1, n])$ computes λ -normal solutions to $\mathcal{T}_2$, enforcing the correct value for λ at every time point using assumptions. The

solutions to the original problem $\mathcal{T}_1$ can be extracted from the λ -normal solutions, after deleting the atoms in $\{\lambda\}[1, n]$.

We turn now our attention to the resolvents δ of the set of nogoods $\Psi_{tr^\lambda(\Pi)}[1, n]$ used by the procedure *CDNL-ASP*. As we will see, just by looking at these resolvents δ , we can approximate the specific interval $[i, j] \subseteq [1, n]$ of the nogoods that were used to prove them.

To this end, we say that the nogoods containing literals of different steps are dynamic nogoods, and they are static nogoods otherwise. All dynamic nogoods in $\Psi_{tr^\lambda(\Pi)}[1, n]$ come from the instantiation of some dynamic integrity constraint $\{\perp \leftarrow B(r) \cup \{\lambda\} \mid r \in \Pi^{di}\}$ at some time step i and, therefore, they contain some literal of the form $\mathbf{T}\lambda_i$. On the other hand, in $\Psi_{tr^\lambda(\Pi)}[1, n]$ there are no literals of the form $\mathbf{F}\lambda_i$. Hence, the literals $\mathbf{T}\lambda_i$ occurring in the dynamic nogoods can never be resolved away. Then, if some dynamic nogood is used to prove a learned nogood δ , the literal $\mathbf{T}\lambda_i$ occurring in that dynamic nogood must belong to δ . This means that the literals $\mathbf{T}\lambda_i$ from a learned nogood δ tell us exactly the steps i of the dynamic nogoods that have been used to prove δ .

Observe now that two nogoods $\delta_1 \in \Psi_{tr^\lambda(\Pi)}[i]$ and $\delta_2 \in \Psi_{tr^\lambda(\Pi)}[i+1]$ can only be resolved if δ_2 is a dynamic nogood. Otherwise, the nogoods would have no opposite literals to resolve. Applying the same reasoning, if two nogoods $\delta_1 \in \Psi_{tr^\lambda(\Pi)}[i]$ and $\delta_2 \in \Psi_{tr^\lambda(\Pi)}[j]$, such that $i < j$, are part of the same resolution proof of a learned nogood δ , then the proof must also contain some dynamic nogoods from each step in the interval $[i+1, j]$. Therefore, the learned nogood δ must contain the literals $\{\mathbf{T}\lambda\}[i+1, j]$.

This implies that, given the literals $\{\mathbf{T}\lambda\}[k, j]$ occurring in a learned nogood δ , we can infer the following about the nogoods from $\Psi_{tr^\lambda(\Pi)}[1, n]$ used to prove δ : dynamic nogoods from all the steps $[k, j]$ were used to prove δ , possibly some static nogoods of the step $k-1$ were used as well, and no nogoods from other steps were used in the proof. It is possible that some static nogoods at steps $[k, j]$ were also used, but no dynamic nogoods at $k-1$ could be used, since otherwise δ should contain the literal $\mathbf{T}\lambda_{k-1}$.

We formalize this with the function $step^\lambda(\delta)$, that approximates the specific interval $[i, j]$ of the nogoods that were used to prove δ : if δ contains some literal of the form $\mathbf{T}\lambda_i$ for $i \in [1, n]$, then $step^\lambda(\delta)$ is the set of steps $\{j-1, j \mid \mathbf{T}\lambda_j \in \delta\}$. For example, if δ is $\{\mathbf{T}a_3, \mathbf{T}\lambda_3\}$ then the value of $step^\lambda(\delta)$ is $\{2, 3\}$. It is clear that δ was derived using some dynamic nogood of step 3, that added the literal $\mathbf{T}\lambda_3$. And it could also happen that some static nogood of step 2 was used, but we are uncertain about it. That is why we say that $step$ is an approximation. To continue, note that it can also be that δ has no literals of the form $\mathbf{T}\lambda_i$. In this case, δ must be the result of resolving some static nogoods of a single time step, and we can extract that time step from the unique time step of the literals occurring in the nogood. Hence, in this case we define $step^\lambda(\delta)$ as $step(\delta)$. For example, $step^\lambda(\{\mathbf{T}c_2, \mathbf{T}d_2\}) = \{2\}$. With this, we can generalize a nogood δ to the shifted nogoods $\delta\langle t \rangle$ whose $step$ value fits in the interval $[1, n]$. We state this precisely in part (i) of the next theorem.

Theorem 4

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem in PNF, and δ be a resolvent of $\Psi_{tr^\lambda(\Pi)}[1, m]$ for some $m \geq 1$. Then, for every $n \geq 1$, the set of nogoods $\Psi_{tr^\lambda(\Pi)}[1, n]$ entails the generalization

$$\{\delta\langle t \rangle \mid step^\lambda(\delta\langle t \rangle) \subseteq [1, n]\}.$$

Observe that part (i) excludes the shifted nogoods $\delta\langle t \rangle$ that contain the literal $\mathbf{T}\lambda_1$, since in that




Fig. 2. Representation of different shifted versions of the nogood $\delta = \{Ta_3, T\lambda_3, T\lambda_4\}$. The surrounding rectangles cover the interval of their *step* value. For example, the rectangle of $\{Ta_2, T\lambda_2, T\lambda_3\}$ covers the interval $[1, 3]$ because $step(\{Ta_2, T\lambda_2, T\lambda_3\}) = [1, 3]$.

case $step^\lambda(\delta\langle t \rangle)$ contains the step $0 \notin [1, n]$. This makes sense because to prove $\delta\langle t \rangle$ we could need some static nogoods at step 0, and they do not belong to $\Psi_{tr^\lambda(\Pi)}[1, n]$.

Example 7

Consider the call $CDNL-ASP(gen(tr^\lambda(\Pi_1), 4), \{\})$, similar to the one that we have seen before using the original program Π_1 . The nogoods $\Psi_{tr^\lambda(\Pi_1)}[1, n]$ are the same as those in $\Psi_{\Pi_1}[1, n]$, except that every dynamic nogood contains one instantiation of the literal $T\lambda$. Instead of learning the nogood $\{Ta_3\}$ (written before as $\{Ta\}[3]$) the algorithm would learn the nogood $\delta = \{Ta_3, T\lambda_3, T\lambda_4\}$. Then, applying part (i) of Theorem 4 the nogood δ can be generalized to $\delta\langle -1 \rangle = \{Ta_2, T\lambda_2, T\lambda_3\}$, but not to $\delta\langle 1 \rangle = \{Ta_4, T\lambda_4, T\lambda_5\}$ or to $\delta\langle -2 \rangle = \{Ta_1, T\lambda_1, T\lambda_2\}$ (see Figure 2).

The next step is to show how temporal programs in general can be translated to PNF form. For this, given a temporal logic program Π over $\mathcal{P}$, let $\mathcal{P}^* = \{p^* \mid p \in \mathcal{P}\}$, and assume that this set is disjoint from $\mathcal{P}$ and $\mathcal{P}'$. The translation $tr^*(\Pi)$ consists of two parts. The first part consists of the result of replacing in Π every atom $p' \in \mathcal{P}'$ by its corresponding new atom p^* . The second part consists of the union of the rules

$$\{\{p^*\} \leftarrow; \perp \leftarrow p', \neg p^*; \perp \leftarrow \neg p', p^*\}$$

for every $p \in \mathcal{P}$. The idea of the translation is that the atoms $p' \in \mathcal{P}'$ are confined to integrity constraints by replacing them by new atoms $p^* \in \mathcal{P}^*$, whose truth value is completely determined by the corresponding $p' \in \mathcal{P}'$ atoms by means of the last set of rules.

Proposition 3

For any temporal logic program Π , the program $tr^*(\Pi)$ is in PNF.

The solutions to temporal problems with Π are the same as the solutions to the same temporal problems with $tr^*(\Pi)$ where the atoms $p^*[i]$ are false at $i = 0$ and have the truth value of $p[i - 1]$ at the other time steps i . Just like before, when we use this translation, we have to add to $CDNL-ASP$ the correct assumptions to fix the value of the p^* atoms at step 0.

Proposition 4

Let $\mathcal{T}_1 = \langle \Pi, I, F \rangle$ and let $\mathcal{T}_2 = \langle tr^*(\Pi), I, F \rangle$ be temporal logic problems. There is a one-to-one correspondence between the solutions to $\mathcal{T}_1$ and the solutions to $\mathcal{T}_2$ that do not contain any atom $p^* \in \mathcal{P}^*$ at step 0.

This proposition allows us to replace any temporal program Π by a temporal program $tr^*(\Pi)$ in PNF. We can then apply the translation tr^λ and benefit from Theorem 4. In fact, we can go one step further, and apply the nogoods learned with the program $tr^\lambda(tr^*(\Pi))$ directly to the original problem with Π . We make this claim precise in the next theorem. For that, we define the simplification of a nogood δ , written $simp(\delta)$, as the nogood $\{\mathbf{V}p_i \mid \mathbf{V}p_i \in \delta, \mathbf{V} \in \{\mathbf{T}, \mathbf{F}\}, p \in \mathcal{P}\} \cup \{\mathbf{V}p_{i-1} \mid \mathbf{V}p_i^* \in \delta, \mathbf{V} \in \{\mathbf{T}, \mathbf{F}\}, p^* \in \mathcal{P}^*\}$ that results from skipping the λ_i literals of δ , and replacing the atoms p_i^* by their corresponding atoms p_{i-1} .

Theorem 5

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem, and δ be a resolvent of $\Psi_{tr^\lambda(tr^*(\Pi))}[1, m]$ for some $m \geq 1$. Then, for every $n \geq 1$, the set of nogoods $\Psi_\Pi[1, n]$ entails the generalization

$$\{simp(\delta\langle t \rangle) \mid step^\lambda(\delta\langle t \rangle) \subseteq [1, n]\}.$$

6 When can we generalize all learned nogoods to all time steps?

Given *any temporal problem*, Theorem 3 gives us a sufficient condition for the generalization of the nogoods learned while solving that problem. In this section, we investigate for *what kind of temporal problems* can we generalize all learned nogoods to all time steps. In other words, we would like to know when can we add the generalization

$$\{\delta\langle t \rangle \mid step(\delta\langle t \rangle) \subseteq [0, n]\}$$

of a learned nogood δ to the set of nogoods used by algorithm *CDNL-ASP*.

Example 8

In Example 6, we saw that the nogood $\delta = \{\mathbf{T}a\}[3]$ is a resolvent of $\Psi_{\Pi_1}[2, 4]$. By Theorem 3, we can conclude that the generalization $\{\{\mathbf{T}a\}[2], \{\mathbf{T}a\}[3]\}$ of δ is entailed by $\Psi_{\Pi_1}[1, n]$. On the other hand, that theorem does not allow us to infer that $\delta\langle -2 \rangle = \{\mathbf{T}a\}[1]$ is entailed by $\Psi_{\Pi_1}[1, n]$. In fact, this would be incorrect since all the solutions to $\Psi_{\Pi_1}[1, n]$ contain the literal $\mathbf{T}a_1$. But why is $\delta\langle -2 \rangle$ not entailed by $\Psi_{\Pi_1}[1, n]$? One reason for this is that $\Psi_{\Pi_1}[1, n]$ does not entail the nogood $\{\mathbf{T}c, \mathbf{F}d\}[0]$, that would be necessary to derive $\delta\langle -2 \rangle$. There are solutions to $\Psi_{\Pi_1}[1, n]$ that violate the nogood $\{\mathbf{T}c, \mathbf{F}d\}[0]$. In fact, since the initial state of all solutions to $\langle \Pi_1, \emptyset, \emptyset \rangle$ is $\{a, b, c\}$, all solutions to $\Psi_{\Pi_1}[1, n]$ contain the literals $\mathbf{T}c_0$ and $\mathbf{F}d_0$, and they all violate that nogood. But let us stop here for a moment, and observe that this implies that the initial state $\{a, b, c\}$ cannot be the next state of some previous state. This is obvious looking at Π_1 , since $\{a, b, c\}$ violates the integrity constraint $\perp \leftarrow c, \neg d$. From the other side, observe that if $\{a, b, c\}$ was the next state of some previous state, then it could not violate that integrity constraint, and therefore $\Psi_{\Pi_1}[1, n]$ would entail $\{\mathbf{T}c, \mathbf{F}d\}[0]$ and then $\delta\langle -2 \rangle$.

This analysis suggests that we can always add a learned nogood, shifted a negative amount of steps, if the initial states have enough previous states. A similar analysis considering the nogood $\delta\langle 1 \rangle$ suggests that we can add a nogood shifted a positive amount of steps if the final states have enough next states. Both suggestions together provide an answer to our question: we can generalize all learned nogoods to all time steps if the initial states have enough previous states, and the final states have enough next states. We make this claim precise in the following.

In Section 3 we introduced transition programs and used them to characterize the solutions of a given temporal problem. Transition programs $trans(\Pi)$ define transitions between the states of some temporal program Π . In turn, these transitions implicitly define a *transition graph* $G(\Pi)$.




Fig. 3. Transition $G(\Pi_1)$ of temporal program Π_1 . The nodes that belong to some solution of length 4 have a gray background. The transitions of those solutions are represented by normal arrows, while the other arrows are dashed.

Formally, given a temporal logic program Π over $\mathcal{P}$, the transition graph $G(\Pi)$ is the graph $\langle N, E \rangle$ where E is the set of edges

$$\{ \langle \{p \mid p' \in X \cap \mathcal{P}'\}, X \cap \mathcal{P} \rangle \mid X \text{ is a stable model of } \text{trans}(\Pi) \}$$

and N is the set of nodes occurring in some edge of E , i.e., $N = \bigcup_{(X,Y) \in E} \{X, Y\}$. With this, we can characterize the solutions to a temporal problem $\langle \Pi, I, F \rangle$ as the finite paths of $G(\Pi)$ whose first and final nodes are consistent with I and F , respectively. We state this formally in the next theorem, that extends Theorem 2.

Theorem 6

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem, n be an integer such that $n \geq 1$, and X be a set of atoms. Then, the following statements are equivalent:

- The pair (X, n) is a solution to $\langle \Pi, I, F \rangle$.
- $X = A^T \cap \mathcal{P}[0, n]$ for a solution A for $\Psi_\Pi[1, n]$ such that $I[0] \cup F[n] \subseteq A$.
- There is a path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ such that $\bigcup_{i \in \{0, \dots, n\}} s_i[i] = X$, the state s_0 is consistent with I , and the state s_n is consistent with F .

Figure 3 represents the transition $G(\Pi_1)$ of temporal program Π_1 . There are only three paths in $G(\Pi)$ of length 4:

- $\langle \{a, b, c\}, \{a, b\}, \{b, d\}, \{c, d\}, \{a, c, d\} \rangle$,
- $\langle \{a, b, c\}, \{a, b\}, \{b\}, \{c, d\}, \{a, c, d\} \rangle$, and
- $\langle \{a, b, c\}, \{a, b\}, \{b\}, \{b, c, d\}, \{a, c, d\} \rangle$.

By Theorem 6, each of them corresponds to one of the solutions to $\langle \Pi_1, \emptyset, \emptyset \rangle$ of length 4.

Theorem 6 establishes a relation between the solutions to the set of nogoods $\Psi_\Pi[1, n]$ and the paths of $G(\Pi)$. This leads naturally to a relation between the nogoods entailed by $\Psi_\Pi[1, n]$ and the paths of $G(\Pi)$.

Example 9

We have seen that the nogood $\delta = \{Ta\}[3]$ is entailed by the set of nogoods $\Psi_{\Pi_1}[2, 4]$. But, what does this mean in $G(\Pi_1)$? If $\delta = \{Ta\}[3]$ is entailed by $\Psi_{\Pi_1}[2, 4]$, then we know that $\delta(-1) = \{Ta\}[2]$ is entailed by $\Psi_{\Pi_1}[1, 3]$. By Theorem 6, when I and F are empty, this implies

that the solutions to $\Psi_{\Pi_1}[1, 3]$ correspond to the paths $\langle s_0, s_1, s_2, s_3 \rangle$ of length 3 in $G(\Pi_1)$. But then, if every solution to $\Psi_{\Pi_1}[1, 3]$ does not violate the nogood $\{\mathbf{T}a\}[2]$, it also has to be the case that the paths $\langle s_0, s_1, s_2, s_3 \rangle$ in $G(\Pi_1)$ do not violate that nogood (we make precise this relation between nogoods and paths below). This gives us directly the answer to our question: if $\Psi_{\Pi_1}[2, 4]$ entails $\{\mathbf{T}a\}[3]$, this means that, in every path $\langle s_0, s_1, s_2, s_3 \rangle$ of $G(\Pi_1)$, the state s_2 cannot contain the atom a . We can check that this is the case in $G(\Pi_1)$, where the only states that appear in such a position are $\{b, d\}$, $\{c, d\}$, $\{b\}$ and $\{b, c, d\}$, and they do not contain the atom a . On the other hand, the states $\{a, b, c\}$ and $\{a, b\}$ can contain a because they are not in the third position of any path, and the same holds for the final state $\{a, c, d\}$, since it does not occur in the penultimate position of any path.

We formalize the relation between nogoods and paths as follows. For simplicity, we only discuss the case where learned nogoods consist of normal atoms, but the extension to body atoms does not pose any special challenge, since body atoms can be seen as a conjunction of normal atoms. Let Π be a temporal program over $\mathcal{P}$, let $\langle s_0, \dots, s_n \rangle$ be some path in $G(\Pi)$, and δ be some (non-temporal) nogood over $\mathcal{P}[0, n]$. We say that the path $\langle s_0, \dots, s_n \rangle$ violates δ if

$$\delta \subseteq \bigcup_{i \in \{0, \dots, n\}} (\{\mathbf{T}p_i \mid p \in s_i\} \cup \{\mathbf{F}p_i \mid p \in \mathcal{P} \setminus s_i\}).$$

The right-hand-side of the equation represents the path as an assignment.

Proposition 5

Let Π be a temporal logic program over $\mathcal{P}$, δ be a (non-temporal) nogood over $\mathcal{P}[k, l]$ for some integers k and l such that $0 \leq k \leq l$, and i and j be some integers such that $i \leq k + 1$ and $l \leq j$. Then, the following two statements are equivalent:

1. The set of nogoods $\Psi_{\Pi}[i, j]$ entails δ .
2. Every path $\langle s_{i-1}, \dots, s_j \rangle$ of length $j + (1 - i)$ in $G(\Pi)$ does not violate δ .

Note that we let $i \leq k + 1$ because $\Psi_{\Pi}[i, j]$ can entail some nogood that contains some atom at step $k = i - 1$. The proof relies on Theorem 6, like we did in Example 9.

Our next proposition states that we can generalize all learned nogoods to all time steps whenever the temporal program is *cyclic*. We say that a temporal program Π is *cyclic* if every node in $G(\Pi)$ belongs to a loop in $G(\Pi)$. In terms of transitions between states, the program Π is cyclic if for any state that participates in some transition there is a sequence of transitions that starts and ends at that state. Observe that if Π is cyclic, then for every path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ and every integers $k \leq 0$ and $l \geq 0$ there is also some path $\langle s_k, \dots, s_0, \dots, s_n, \dots, s_{n+l} \rangle$ that extends the original path from both ends. This property is crucial for the generalization to all time steps. It implies that whenever a path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ does not violate a nogood δ over $\mathcal{P}[i, j]$ where $0 \leq i \leq j \leq n$, then every (possibly smaller) path $\langle s_i, \dots, s_j \rangle$ in $G(\Pi)$ also does not violate δ . Consider the case where a set of nogoods $\Psi_{\Pi}[1, n]$ entails δ . By Proposition 5 we know that every path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ does not violate δ . Then, this property of cyclic programs allows us to infer that every (possibly smaller) path $\langle s_i, \dots, s_j \rangle$ in $G(\Pi)$ also does not violate δ , and then Proposition 5 again gives us that $\Psi_{\Pi}[i + 1, j]$ entails δ . Since the bounds that we obtain, $i + 1$ and j , are tight as possible for the learned nogood δ , now we can shift δ to all possible time steps inside $[0, n]$ and the resulting nogood is always entailed by $\Psi_{\Pi}[1, n]$.

Example 10

Consider a temporal logic program Π_2 that is cyclic, together with the set of nogoods $\Psi_{\Pi_2}[1, 6]$, and a nogood $\{\mathbf{T}a_3, \mathbf{T}b_4\}$ that is a resolvent of $\Psi_{\Pi_2}[2, 5]$. By Theorem 3, we can conclude that both $\{\mathbf{T}a_2, \mathbf{T}b_3\}$ and $\{\mathbf{T}a_4, \mathbf{T}b_5\}$ are entailed by $\Psi_{\Pi_2}[1, 6]$, but we cannot do the same about $\{\mathbf{T}a_0, \mathbf{T}b_1\}$, $\{\mathbf{T}a_1, \mathbf{T}b_2\}$, or $\{\mathbf{T}a_5, \mathbf{T}b_6\}$. On the other hand, by Proposition 5, the fact that $\Psi_{\Pi_2}[1, 6]$ entails $\{\mathbf{T}a_3, \mathbf{T}b_4\}$ implies that for every path $\langle s_0, s_1, s_2, s_3, s_4, s_5, s_6 \rangle$ in $G(\Pi)$ it cannot be the case that s_3 contains a and s_4 contains b . Then, given that Π is cyclic, it holds that for every path $\langle s_3, s_4 \rangle$ in $G(\Pi)$ it cannot be the case that s_3 contains a and s_4 contains b . To see this, observe that otherwise there would be some path $\langle s_0, s_1, s_2, s_3, s_4, s_5, s_6 \rangle$ in $G(\Pi)$ that would violate those conditions on s_3 and s_4 , which would contradict our previous statement. But then, it holds that the paths $\langle s_3, s_4 \rangle$ in $G(\Pi)$ do not violate the nogood $\{\mathbf{T}a_3, \mathbf{T}b_4\}$. In this case, Proposition 5 tells us that the original nogood $\{\mathbf{T}a_3, \mathbf{T}b_4\}$ is entailed by $\Psi_{\Pi_2}[4, 4]$, while before we only knew that it was entailed by $\Psi_{\Pi_2}[2, 5]$. From this, it follows that $\Psi_{\Pi_2}[1, 1]$ entails the nogood $\{\mathbf{T}a_0, \mathbf{T}b_1\}$, $\Psi_{\Pi_2}[2, 2]$ entails the nogood $\{\mathbf{T}a_1, \mathbf{T}b_2\}$, and $\Psi_{\Pi_2}[6, 6]$ entails the nogood $\{\mathbf{T}a_5, \mathbf{T}b_6\}$. Then, the set of nogoods $\Psi_{\Pi_2}[1, 6]$ also entails those three nogoods, and in total it entails all shifted versions of $\{\mathbf{T}a_3, \mathbf{T}b_4\}$ that fit in the interval $[0, 6]$.

Theorem 7

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem, and δ be a resolvent of $\Psi_{\Pi}[i, j]$ for some i and j such that $1 \leq i \leq j$. If Π is cyclic, for every $n \geq 1$, the set of nogoods $\Psi_{\Pi}[1, n]$ entails the generalization

$$\{\delta\langle t \rangle \mid \text{step}(\delta\langle t \rangle) \subseteq [0, n]\}.$$

Theorem 7 allows us to generalize all learned nogoods to all time steps whenever a temporal program is cyclic. Given that we are specially interested in planning problems, the question for us is what kind of planning problems can be represented using these cyclic temporal programs. One example of this class are planning problems that satisfy the following two conditions:

1. Whenever no action occurs in a state, that state remains the same.
2. Whenever an action occurs in a state, the repetition of that action in that state does not change the state.

The first condition creates a one-step loop in all states without action occurrences, and the second does the same in all states with action occurrences. In this way, all states, with or without action occurrences, belong to a loop, and therefore the temporal program is cyclic.

Certainly, there are problems where the temporal programs are cyclic, but there are many others that do not fall into this class, and we would like to increase the applicability of our approach. For this reason, we are going to extend our study and also take into account the initial states of a temporal logic problem. Once we do this, we will no longer require the generalization of the learned nogoods to be entailed by a set of nogoods, but it will be enough for us if the generalization does not violate any solution that is consistent with the initial situation. This condition is sufficient to allow us to add the generalized nogoods to the *CDNL-ASP* algorithm.

Let $\mathcal{T} = \langle \Pi, I, F \rangle$ be a temporal problem. The *initial states* of $\mathcal{T}$ are the states of $G(\Pi)$ that are consistent with I . The *reachable states* of $\mathcal{T}$ are the states of $G(\Pi)$ that are reachable in $G(\Pi)$ from the initial states of $\mathcal{T}$. Note that this includes the initial states of $\mathcal{T}$. The temporal problem $\mathcal{T}$ is *cyclic* if these conditions hold:

1. Every initial state of $\mathcal{T}$ belongs to a loop in $G(\Pi)$.

2. Every reachable state s of $\mathcal{T}$ is connected by an edge in $G(\Pi)$ to some state s' that belongs to a loop in $G(\Pi)$.

Cyclic temporal problems have a similar property like the one for cyclic temporal programs: if $\mathcal{T}$ is cyclic, then for every path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ such that s_0 is consistent with I and every integers $k \leq 0$ and $l \geq 0$, there is some path $\langle s_k, \dots, s_0, \dots, s_n, \dots, s_{n+l} \rangle$ in $G(\Pi)$ that extends the original path from both ends. The states $s_k, \dots, s_0$ may belong to some loop, that must exist by the previous first condition on cyclic problems, and the states $s_{n+1}, \dots, s_{n+l}$ may belong to another loop, that must exist by the previous second condition.

Theorem 8

Let $\langle \Pi, I, F \rangle$ be a temporal logic problem, and δ be a resolvent of $\Psi_\Pi[i, j]$ for some i and j such that $1 \leq i \leq j$. If $\langle \Pi, I, F \rangle$ is cyclic, for every $n \geq 1$, the nogoods that belong to the generalization

$$\{\delta\langle t \rangle \mid \text{step}(\delta\langle t \rangle) \subseteq [0, n]\}$$

are not violated by any solution A for $\Psi_\Pi[1, n]$ that is consistent with $I[0]$.

Theorem 8 allows us to generalize all learned nogoods to all time steps whenever the temporal problem is cyclic. Now, it turns out that most planning problems that we have found fall into this cyclic class. This is the case whenever they satisfy these conditions:

1. The initial state has no action occurrences.
2. Whenever no action occurs in a state, that state remains the same.
3. There is a transition from every state with some action occurrence to the same state without action occurrences.

Observe how the first and the second condition imply the first condition of cyclic problems, while the second and the third condition imply the second condition of cyclic problems.

It turns out that planning problems usually satisfy the three conditions mentioned above. To satisfy the first one it is enough to place the action occurrences in the state where their effects take place. Usually, in ASP, this is a matter of convenience. To satisfy the second and the third conditions it is enough if the representation allows for the non-execution of actions at every state and inertia makes everything persist. Usually, planning problem descriptions allow for that, or they can be easily modified to do that. In practice, the encodings that we used in our conference paper (Romero et al. 2022) only required minor changes to fit into this class.

7 Experiments

In this section, we experimentally evaluate the generalization of learned nogoods in ASP planning using the solver *clingo*. The goal of the experiments is to study the performance of *clingo* when the planning encodings are extended by the generalizations of some constraints learned by *clingo* itself. We are interested only in the solving time and not in the grounding time, but in any case we have observed no differences between grounding times among the different configurations compared. We performed experiments in two different settings, single shot and multi shot, that we detail below. Following the approach of (Gebser et al. 2016), in all experiments we disregarded the learned nogoods of size greater than 50 and of degree greater than 10, where the degree of a nogood is defined as the difference between the maximum and minimum step of the literals of the nogood. In all the experiments, the learned nogoods are always sorted either by size or by

literal block distance (*lbd*, (Audemard and Simon 2009)), a measure that is usually associated with the quality of a learned nogood. We tried configurations adding the best 500, 1000, or 1500 nogoods, according to either their nogood size or their *lbd*. The results ordering the nogoods by *lbd* were similar but slightly better than those ordering by size, and here we focus on them. We used two benchmark sets from (Dimopoulos et al. 2018). The first consists of PDDL benchmarks from planning competitions, translated to ASP using the system *plasp* presented in that paper. This set contains 120 instances of 6 different domains. The second set consists of ASP planning benchmarks from ASP competitions. It contains 136 instances of 9 domains. We adapted the logic programs of these benchmarks to the format of temporal logic programs as follows: we deleted the facts used to specify the initial situation, as well as the integrity constraints used to specify the goal, we added some choice rules to open the initial situation, and we fixed the initial situation and the goal using assumptions. All benchmarks were run using the version 5.5.1 of *clingo* on an Intel Xeon E5-2650v4 under Debian GNU/Linux 10, with a memory limit of 8 GB, and a timeout of 15 minutes per instance.

The task in the single shot experiment is to find a plan of a fixed length n that is part of the input. For the PDDL benchmarks we consider plan lengths varying from 5 to 75 in steps of 5 units, for a total of 2040 instances. The ASP benchmarks already have a plan length, and we use it. In a preliminary learning step, *clingo* is run with every instance for 10 minutes or until 16000 nogoods are learned, whatever happens first. The actual learning time is disregarded and not taken into account in the tables. Some PDDL instances overcome the memory limit in this phase. We leave them aside and are left with 1663 instances of this type. We compare the performance of *clingo* running normally (baseline), versus the (learning) configurations where we add the best 500, 1000, or 1500 learned nogoods according to their *lbd* value. In this case we apply Theorem 8 and learn the nogoods using a slight variation of the original encoding, but use the original encoding for the evaluation of all configurations. The results applying the translations from Section 5 are similar. They can be found in the Appendix Appendix B.

Tables 1 and 2 show the results for the PDDL and the ASP benchmarks, respectively. The first columns include the name and number of instances of every domain. The tables show the average solving times and the number of timeouts, in parenthesis, for every configuration and domain. We can observe that in general the learning configurations are faster than the baseline and in some domains they solve more instances. The improvement is not huge, but is persistent among the different settings. The only exception is the *elevator* domain in PDDL, where the baseline is a bit faster than the other configurations. We also analyzed the average number of conflicts per domain and configuration, and the results follow the same trend as the solving times.

		baseline	500	1000	1500
blocks	(300)	0.5 (0)	0.6 (0)	0.6 (0)	0.6 (0)
depots	(270)	146.8 (30)	140.2 (30)	124.2 (24)	135.5 (28)
driverlog	(135)	14.0 (1)	13.5 (1)	11.5 (1)	10.8 (1)
elevator	(300)	3.0 (0)	5.1 (0)	4.3 (0)	5.2 (0)
grid	(30)	11.4 (0)	6.0 (0)	4.4 (0)	3.7 (0)
gripper	(255)	381.1 (96)	380.9 (90)	360.9 (87)	367.7 (90)
logistics	(225)	0.5 (0)	0.5 (0)	0.5 (0)	0.8 (0)
mystery	(130)	79.6 (6)	71.1 (3)	58.6 (4)	64.6 (6)
Total	(1645)	91.5 (133)	90.0 (124)	83.0 (116)	86.5 (125)

Table 1. Single shot solving of PDDL benchmarks.

		baseline	500	1000	1500
HanoiTower	(20)	160.5 (2)	139.5 (0)	137.9 (0)	143.9 (1)
Labyrinth	(20)	246.6 (3)	348.3 (5)	284.8 (4)	296.2 (5)
Nomistery	(20)	585.7 (12)	545.4 (11)	510.2 (9)	566.7 (12)
Ricochet Robots	(20)	464.5 (9)	320.3 (3)	410.8 (6)	404.9 (5)
Sokoban	(20)	458.7 (9)	454.2 (9)	453.8 (9)	446.3 (9)
Visit-all	(20)	559.1 (12)	562.5 (12)	560.7 (12)	561.5 (12)
Total	(120)	412.5 (47)	395.0 (40)	393.0 (40)	403.3 (44)

Table 2. Single shot solving of ASP benchmarks.

In the Multi shot solving experiment, the solver first looks for a plan of length 5. If the solver returns that there is no such plan, then it looks for a plan of length 10, and so on until it finds a plan. At each of these solver calls, we collect the best learned nogoods. Then, before the next solver call, we add the generalization of the best 500, 1000, or 1500 of them, depending on the configuration. As before, we rely on Theorem 8, but this time we use the same original encoding, slightly modified, for both learning and solving. The results using the translations from Section 5 are similar. They can be found in the Appendix Appendix B.

The results for PDDL and ASP are shown in Tables 3 and 4, respectively. In both of them, the baseline and the different configurations perform similarly, and we do not observe a clear trend. The analysis of the average number of conflicts shows similar results.

We expected similar results on the single shot and the multi shot solving experiments. However, this is not what we have observed. The learning configurations perform better than the baseline in the former, but similarly in the latter. At the moment, we have found no clear explanation for these results, but we hope to make some progress in this regard in the future.

		baseline	500	1000	1500
blocks	(20)	1.3 (0)	0.7 (0)	0.7 (0)	0.7 (0)
depots	(18)	148.6 (2)	255.9 (3)	188.7 (3)	221.9 (3)
driverlog	(9)	108.8 (1)	102.1 (1)	104.8 (1)	108.6 (1)
elevator	(20)	280.4 (5)	285.6 (5)	293.8 (5)	304.6 (5)
freecell	(16)	900.0 (16)	900.0 (16)	900.0 (16)	900.0 (16)
grid	(2)	5.1 (0)	3.9 (0)	4.1 (0)	4.3 (0)
gripper	(17)	848.6 (16)	847.5 (16)	849.0 (16)	847.9 (16)
logistics	(20)	225.2 (5)	225.3 (5)	225.4 (5)	225.3 (5)
mystery	(14)	321.8 (5)	321.9 (5)	321.9 (5)	321.9 (5)
Total	(136)	346.6 (50)	360.9 (51)	353.6 (51)	359.7 (51)

Table 3. Multi shot solving of PDDL benchmarks.

		baseline	500	1000	1500
HanoiTower	(20)	440.8 (8)	512.8 (9)	489.4 (9)	498.9 (9)
Labyrinth	(20)	633.9 (14)	633.8 (14)	633.8 (14)	633.9 (14)
Nomistery	(20)	380.7 (7)	363.1 (6)	381.0 (7)	384.7 (7)
Ricochet Robots	(20)	521.5 (11)	523.9 (11)	527.9 (11)	526.0 (11)
Sokoban	(20)	721.5 (16)	721.5 (16)	721.9 (16)	722.1 (16)
Visit-all	(20)	900.0 (20)	900.0 (20)	900.0 (20)	900.0 (20)
Total	(120)	599.7 (76)	609.2 (76)	609.0 (77)	610.9 (77)

Table 4. Multi shot solving of ASP benchmarks.

8 Conclusion

Conflict-driven constraint learning (CDCL) is the key to the success of modern ASP solvers. So far, however, ASP solvers could not exploit the temporal structure of dynamic problems. We addressed this by elaborating upon the generalization of learned constraints in ASP solving for temporal domains. We started with the definition of temporal logic programs and problems. For temporal programs that only refer to previous time steps within integrity constraints, we show that a simple modification of the temporal program is enough to generalize all learned constraints to all time points. This is no real restriction because any temporal program can be translated into this restricted format. Hence, once we apply both translations, we have a representation whose nogoods can always be generalized. In addition to this, we identified a class of temporal problems for which every learned nogood can be generalized to all time points, without the need of any translation. It turns out that many planning problems fall into this case. Our experiments show mixed results. In some settings, the addition of the learned constraints results in a consistent improvement of performance, while in others the performance is similar to the baseline. We plan to continue this experimental investigation in the future. Another avenue of future work is to continue the approach sketched at the end of Section 4, and develop a dedicated implementation within an ASP solver based in Theorem 3.

Acknowledgments.

This work was supported by DFG grant SCHA 550/15.

Competing interests

The authors declare none.

References

- AGUADO, F., CABALAR, P., DIÉGUEZ, M., PÉREZ, G., AND VIDAL, C. 2013. Temporal equilibrium logic: a survey. *Journal of Applied Non-Classical Logics* 23, 1-2, 2–24.
- AUDEMARD, G. AND SIMON, L. 2009. Predicting learnt clauses quality in modern SAT solvers. In *Proceedings of the Twenty-first International Joint Conference on Artificial Intelligence (IJCAI’09)*, C. Boutilier, Ed. AAAI/MIT Press, 399–404.
- CABALAR, P., DIÉGUEZ, M., HAHN, S., AND SCHAUB, T. 2021. Automata for dynamic answer set solving: Preliminary report. In *Proceedings of the Fourteenth Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP’21)*.
- COMPLOI-TAUPE, R., WEINZIERL, A., AND FRIEDRICH, G. 2020. Conflict generalisation in asp: Learning correct and effective non-ground constraints. *Theory and Practice of Logic Programming* 20, 799–814.
- CUTERI, B., DODARO, C., RICCA, F., AND SCHÜLLER, P. 2020. Overcoming the grounding bottleneck due to constraints in ASP solving: Constraints become propagators. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI’20)*, C. Bessière, Ed. ijcai.org, 1688–1694.
- DIMOPOULOS, Y., GEBSER, M., LÜHNE, P., ROMERO, J., AND SCHAUB, T. 2018. plasp 3: Towards effective ASP planning. *Theory and Practice of Logic Programming* 19, 3, 477–504.
- EÉN, N. AND SÖRENSON, N. 2003. Temporal induction by incremental SAT solving. *Electronic Notes in Theoretical Computer Science* 89, 4.
- FANDINNO, J., LAFERRIERE, F., ROMERO, J., SCHAUB, T., AND SON, T. 2021. Planning with incomplete information in quantified answer set programming. *Theory and Practice of Logic Programming* 21, 5, 663–679.

- GEBSER, M., KAMINSKI, R., KAUFMANN, B., LÜHNE, P., ROMERO, J., AND SCHAUB, T. 2016. Answer set solving with generalized learned constraints. In *Technical Communications of the Thirty-second International Conference on Logic Programming (ICLP'16)*, M. Carro and A. King, Eds. OpenAccess Series in Informatics (OASICS), vol. 52. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 9:1–9:15.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., AND SCHAUB, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool Publishers.
- GEBSER, M., KAUFMANN, B., NEUMANN, A., AND SCHAUB, T. 2007. Conflict-driven answer set solving. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, M. Veloso, Ed. AAAI/MIT Press, 386–392.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proceedings of the Fifth International Conference and Symposium of Logic Programming (ICLP'88)*, R. Kowalski and K. Bowen, Eds. MIT Press, 1070–1080.
- GELFOND, M. AND LIFSCHITZ, V. 1998. Action languages. *Electronic Transactions on Artificial Intelligence* 3, 6, 193–210.
- LEFÈVRE, C., BÉATRIX, C., STÉPHAN, I., AND GARCIA, L. 2017. ASPeRiX, a first-order forward chaining approach for answer set computing. *Theory and Practice of Logic Programming* 17, 3, 266–310.
- LIFSCHITZ, V. AND TURNER, H. 1994. Splitting a logic program. In *Proceedings of the Eleventh International Conference on Logic Programming*. MIT Press, 23–37.
- PALÙ, A. D., DOVIER, A., PONTELLI, E., AND ROSSI, G. 2009. GASP: Answer set programming with lazy grounding. *Fundamenta Informaticae* 96, 3, 297–322.
- RINTANEN, J., HELJANKO, K., AND NIEMELÄ, I. 2006. Planning as satisfiability: parallel plans and algorithms for plan search. *Artificial Intelligence* 170, 12-13, 1031–1080.
- ROMERO, J., SCHAUB, T., AND STRAUCH, K. 2022. On the generalization of learned constraints for ASP solving in temporal domains. In *Rules and Reasoning - 6th International Joint Conference on Rules and Reasoning, RuleML+RR 2022, Berlin, Germany, September 26-28, 2022, Proceedings*, G. Governatori and A. Turhan, Eds. Lecture Notes in Computer Science, vol. 13752. Springer, 20–37.
- WEINZIERL, A., TAUPE, R., AND FRIEDRICH, G. 2020. Advancing lazy-grounding ASP solving techniques — restarts, phase saving, heuristics, and more. *Theory and Practice of Logic Programming* 20, 5, 609–624.

Appendix A Proofs

Lemma 1

For any temporal logic program Π , $\Sigma_{\Pi}[n] = \Sigma_{\Pi[n]}$.

Lemma 1

1. Constraints

For any constraint c of the form $\perp \leftarrow a_1, \dots, a_m, \text{not } b_{m+1}, \dots, \text{not } b_l$ it holds that

$$\begin{aligned} c[n] &= \perp \leftarrow a_1[n], \dots, a_m[n], \text{not } b_{m+1}[n], \dots, \text{not } b_l[n] \\ \Sigma_{c[n]} &= \mathbf{T}a_1[n], \dots, \mathbf{T}a_m[n], \mathbf{F}b_{m+1}[n], \dots, \mathbf{F}b_l[n] \\ \Sigma_c &= \mathbf{T}a_1, \dots, \mathbf{T}a_m, \mathbf{F}b_{m+1}, \dots, \mathbf{F}b_l \\ \Sigma_c[n] &= \mathbf{T}a_1[n], \dots, \mathbf{T}a_m[n], \mathbf{F}b_{m+1}[n], \dots, \mathbf{F}b_l[n] \end{aligned}$$

Since all constraints have the form of c , we can conclude that $\Sigma_{c[n]} = \Sigma_c[n]$ for any constraint.

2. Body

For any body B of the form $\{a_1, \dots, a_m, \text{not } b_{m+1}, \dots, \text{not } b_l\}$ it holds that

$$\begin{aligned} B[n] &= a_1[n], \dots, a_m[n], \text{not } b_{m+1}[n], \dots, \text{not } b_l[n] \\ \Sigma_{B[n]} &= \{\{\mathbf{T}B[n], \mathbf{F}a_1[n]\}, \dots, \{\mathbf{T}B[n], \mathbf{F}b_l[n]\}\} \cup \{\mathbf{F}B[n], \mathbf{T}a_1[n], \dots, \mathbf{F}b_l[n]\} \\ \Sigma_B &= \{\{\mathbf{T}B, \mathbf{F}a_1\}, \dots, \{\mathbf{T}B, \mathbf{F}b_l\}\} \cup \{\mathbf{F}B, \mathbf{T}a_1, \dots, \mathbf{F}b_l\} \\ \Sigma_B[n] &= \{\{\mathbf{T}B[n], \mathbf{F}a_1[n]\}, \dots, \{\mathbf{T}B[n], \mathbf{F}b_l[n]\}\} \cup \{\mathbf{F}B[n], \mathbf{T}a_1[n], \dots, \mathbf{F}b_l[n]\} \end{aligned}$$

Since all bodies have the form of B , we can then conclude that $\Sigma_{B[n]} = \Sigma_B[n]$.

3. A set of rules with the same head

For any set of rules with the same head Π of the form $\{a \leftarrow B_1, \dots, a \leftarrow B_l\}$ where B_i are bodies and Σ_B is the set of all body nogoods it holds that

$$\begin{aligned} \Pi[n] &= \{a[n] \leftarrow B_1[n], \dots, a[n] \leftarrow B_l[n]\} \\ \Sigma_{\Pi[n]} &= \{\{\mathbf{F}B_1[n], \dots, \mathbf{F}B_l[n], \mathbf{T}a[n]\}, \{\mathbf{T}B_1[n], \mathbf{F}a[n]\}, \dots, \{\mathbf{T}B_l[n], \mathbf{F}a[n]\}\} \cup \Sigma_B \\ \Sigma_{\Pi} &= \{\{\mathbf{F}B_1, \dots, \mathbf{F}B_l, \mathbf{T}a\}, \{\mathbf{T}B_1, \mathbf{F}a\}, \dots, \{\mathbf{T}B_l, \mathbf{F}a\} \cup \text{BodyNogoods}\} \\ \Sigma_{\Pi}[n] &= \{\{\mathbf{F}B_1[n], \dots, \mathbf{F}B_l[n], \mathbf{T}a[n]\}, \{\mathbf{T}B_1[n], \mathbf{F}a[n]\}, \dots, \{\mathbf{T}B_l[n], \mathbf{F}a[n]\}\} \cup \Sigma_B \end{aligned}$$

Since Body nogoods are also equal, we can conclude that $\Sigma_{\Pi[n]} = \Sigma_{\Pi}[n]$.

4. Choice rules

Since choice rule nogoods are a subset of normal rule nogoods, we can conclude that $\Sigma_{c[n]} = \Sigma_c[n]$ for any choice rule c .

5. Loops

For any set of rules Π forming a loop of the form $a_1 \leftarrow a_2, B_1, \dots, a_n \leftarrow a_1, B_n$ with external Bodies for (some) a_i being labeled E_i and Σ_R is the set of all rule nogoods it holds that

$$\begin{aligned} \Pi[n] &= \{a_1[n] \leftarrow a_2[n], B_1[n], \dots, a_n[n] \leftarrow a_1[n], B_n[n]\} \\ \Sigma_{\Pi[n]} &= \{\{\mathbf{T}a_1[n], \mathbf{F}E_{i_1}[n], \dots, \mathbf{F}E_{i_m}[n]\}, \dots, \{\mathbf{T}a_n[n], \mathbf{F}E_{i_1}[n], \dots, \mathbf{F}E_{i_m}[n]\}\} \cup \Sigma_R \\ \Sigma_{\Pi} &= \{\{\mathbf{T}a_1, \mathbf{F}E_{i_1}, \dots, \mathbf{F}E_{i_m}\}, \dots, \{\mathbf{T}a_n, \mathbf{F}E_{i_1}, \dots, \mathbf{F}E_{i_m}\}\} \cup \Sigma_R \\ \Sigma_{\Pi}[n] &= \{\{\mathbf{T}a_1[n], \mathbf{F}E_{i_1}[n], \dots, \mathbf{F}E_{i_m}[n]\}, \dots, \{\mathbf{T}a_n[n], \mathbf{F}E_{i_1}[n], \dots, \mathbf{F}E_{i_m}[n]\}\} \cup \Sigma_R \end{aligned}$$

Since rule nogoods are also equal, we can conclude that $\Sigma_{\Pi[n]} = \Sigma_{\Pi}[n]$.

From items 1, 2, 3, 4, 5 we can say that for any program Π , $\Sigma_{\Pi[n]} = \Sigma_{\Pi}[n]$.

□

Proposition 1

Let $C = C = \{\{p'\} \leftarrow |p \in \mathcal{P}\}$ where $\mathcal{P}$ is the set of atoms occurring in Π . Since all the rules in C are choice rules with empty bodies, $\Sigma_{C[n]}$ is comprised of nogoods of the form $\{\mathbf{T}p'[n], \mathbf{F}\emptyset\}$. Given that $\mathbf{F}\emptyset$ is always false the nogoods can be safely removed. Hence, for any program Π it holds that $\Sigma_{C[n]} \cup \Sigma_{\Pi} = \Sigma_{\Pi}$.

For a given temporal logic program Π we can define $\text{trans}(\Pi) = C \cup \Pi$. Additionally, $\text{gen}(\Pi, n)$

can be defined as $C[1] \cup \Pi[1, n]$, which means that

$$\begin{aligned}\Sigma_{gen(\Pi, n)} &= \Sigma_{C[1]} \cup \Sigma_{\Pi[1, n]} \\ &= \Sigma_{C[1]} \cup \Sigma_{\Pi[1]} \cup \dots \cup \Sigma_{\Pi[n]} \\ &= \Sigma_{\Pi[1]} \cup \dots \cup \Sigma_{\Pi[n]} \quad (\text{deleting choice nogoods})\end{aligned}$$

Also,

$$\begin{aligned}\Psi_{\Pi}[1, n] &= \Sigma_{trans(\Pi)}[1, n] \\ &= \Sigma_{trans(\Pi)}[1] \cup \dots \cup \Sigma_{trans(\Pi)}[n] \\ &= \Sigma_C[1] \cup \Sigma_{\Pi}[1] \cup \dots \cup \Sigma_C[n] \cup \Sigma_{\Pi}[n] \\ &= \Sigma_{\Pi}[1] \cup \dots \cup \Sigma_{\Pi}[n] \quad (\text{deleting choice nogoods}) \\ &= \Sigma_{\Pi[1]} \cup \dots \cup \Sigma_{\Pi[n]} \quad (\text{lemma1}) \\ &= \Sigma_{gen(\Pi, n)}\end{aligned}$$

□

Theorem 2

Let $\mathcal{P}$ be the set of atoms occurring in Π .

By proposition 1 a solution for the set of nogoods $\Psi_{\Pi}[1, n]$ is a solution for $\Sigma_{gen(\Pi, n)}$. A solution for $\Sigma_{gen(\Pi, n)}$ is a stable model for the generator program $gen(\Pi, n)$. Since a stable model of $gen(\Pi, n)$ consistent with I and F is a solution of $\langle \Pi, I, F \rangle$, then a solution A of $\Psi_{\Pi}[1, n]$ consistent with some I and F , the pair (X, n) where $X = A^T \cap \mathcal{P}[1, n]$ is a solution for $\langle \Pi, I, F \rangle$.

Let (X, n) be a solution to the temporal logic problem $\langle \Pi, I, F \rangle$. By definition, X is a stable model of $gen(\Pi, n)$ consistent with I and F . Since a stable model of $gen(\Pi, n)$ is a solution of $\Sigma_{gen(\Pi, n)}$ which is a solution of $\Psi_{\Pi}[1, n]$ (by proposition 1), it follows that $A = \{Tp | p \in X\} \cup \{Fp | p \in P[0, n] \setminus X\}$ is a solution for the temporal logic program $\Psi_{\Pi}[1, n]$ such that $I[0] \cup F[n] \subseteq A$. □

Lemma 2

For any resolvent δ of $\Psi[i, j]$ it holds that $\delta\langle t \rangle$ is a resolvent of $\Psi[i + t, j + t]$

Lemma 2

Recall that if a nogood is a resolvent of $\Psi[i, j]$ then it must have a resolution proof $\mathcal{T}$ where every nogood $\delta_i \in \mathcal{T}$ is either entailed by $\Psi[i, j]$ or the result of resolving some δ_j and δ_k where $j < k < i$ and both δ_j and δ_k are entailed by $\Psi[i, j]$. Additionally, for a resolution proof $\mathcal{T} = \delta_1, \dots, \delta_n$ the result is δ_n . Finally, note that if a nogood $\delta \in \Psi[i, j]$ then $\delta\langle t \rangle \in \Psi[i + t, j + t]$

We now prove the lemma by induction. Let $\mathcal{T}$ be the resolution proof of a nogood δ that is entailed by $\Psi[i, j]$.

Induction base 1: If $\mathcal{T} = \delta$ then $\delta \in \Psi[i, j]$ holds and, trivially, $\delta\langle t \rangle \in \Psi[i + t, j + t]$.

Induction base 2: If $\mathcal{T} = \delta_1, \delta_2$ then, since there less than two nogoods before δ_1 and δ_2 then both must be in $\Psi[i, j]$. Consequently, $\delta_1\langle t \rangle \in \Psi[i + t, j + t]$ and $\delta_2\langle t \rangle \in \Psi[i + t, j + t]$.

Induction step n: Let $\mathcal{T} = \delta_1, \dots, \delta_n$ be a resolution proof for nogood δ_n . If $\delta_n \in \Psi[i, j]$ then, trivially, $\delta_n\langle t \rangle \in \Psi[i + t, j + t]$. If $\delta_n \notin \Psi[i, j]$ then we know by induction that all δ_i with $0 \leq i \leq n - 1$ are entailed by $\Psi[i, j]$. Since $\delta_n \notin \Psi[i, j]$ then there are some δ_k and δ_l

where $k < l < i$ that resolve to δ_n . By induction, $\delta_k \langle t \rangle$ and $\delta_l \langle t \rangle$ are entailed by $\Psi[i + t, j + t]$. Consequently, $\delta_n \langle t \rangle$ is entailed by $\Psi[i + t, j + t]$.

□

Theorem 3

Let δ be a resolvent of $\Psi[i, j]$. Then the shifted nogood $\delta \langle t \rangle$ is entailed by $\Psi[i + t, j + t]$ (lemma 2). Let t be a value where $[i + t, j + t] \subseteq [1, n]$ holds, then $\delta \langle t \rangle$ is entailed by $\Psi[1, n]$ since $\Psi[i + t, j + t] \subseteq \Psi[1, n]$.

□

Proposition 2

For any model I of $\mathcal{T}_1$ of length n there is also a model $I^\lambda = I \cup \{\lambda_1, \dots, \lambda_n\}$ of $\mathcal{T}_2$. Since tr^λ only adds a λ to the dynamic constraints the only difference in the nogoods of $\mathcal{T}_1$ and $\mathcal{T}_2$ is that $tr^\lambda(\Pi)^{di}$ have an additional λ . Hence, no nogoods of $tr^\lambda(\Pi) \setminus tr^\lambda(\Pi)^{di}$ is satisfied by I^λ . Additionally, since in λ -normal solutions all λ are true, the nogoods of Π^{di} can be simplified by deleting their λ . The simplified nogoods of $tr^\lambda(\Pi)^{di}$ are the same as the nogoods of Π^{di} . This means that I^λ does not satisfy any nogood in $tr^\lambda(\Pi)^{di}$. We can then conclude that I^λ does not satisfy any nogood of $tr^\lambda(\Pi)$ and is thus a model of $\mathcal{T}_2$.

For any model I^λ of $\mathcal{T}_2$ of length n there is also a model $I = I^\lambda \setminus \{\lambda_1, \dots, \lambda_n\}$ of $\mathcal{T}_1$. Since tr^λ only adds a λ to the dynamic constraints the only difference in the nogoods of $\mathcal{T}_1$ and $\mathcal{T}_2$ is that $tr^\lambda(\Pi)^{di}$ have an additional λ . Hence, all nogoods of $\Pi \setminus \Pi^{di}$ are not satisfied by I . Since in λ -normal solutions all λ are true, the nogoods of $tr^\lambda(\Pi)^{di}$ act the same way as the nogoods of Π^{di} . Hence, the nogoods of Π^{di} are also not satisfied by I . We can then conclude that I does not satisfy any nogood of Π and is thus a model of $\mathcal{T}_1$.

It follows that that for every stable model of $\langle \Pi, I, F \rangle$ there is a corresponding stable model of $\langle tr^\lambda(\Pi), I, F \rangle$ and vice versa

□

Lemma 3

For some resolvent δ of $\Psi_{tr}^\lambda(\Pi)[i, j]$, $step^\lambda(\delta)$ computes an *overapproximation* of the interval $[i, j]$.

Lemma 3

- **case 1:** δ is a resolvent of $\Psi_{tr}^\lambda(\Pi)[i, i]$. By definition, $step^\lambda(\delta) = [i, i]$ since there would be no λ in delta and the only timestep in atoms of δ would be i .
- **case 2:** δ is a resolvent of $\Psi_{tr}^\lambda(\Pi)[i, j]$ with $i < j$ and $\lambda[i, j] \in \delta$. By definition, $step^\lambda(\delta) = [i - 1, j]$ since the lowest timestep in any λ is i .
- **case 3:** δ is a resolvent of $\Psi_{tr}^\lambda(\Pi)[i, j]$ with $i < j$ and $\lambda[i + 1, j] \in \delta$. By definition, $step^\lambda(\delta) = [i, j]$ since the lowest timestep in any λ is $i + 1$.

We can clearly see that for any resolvent δ the function $step^\lambda(\delta)$ computes the exact (cases 1 and 3) or a bigger (case 2) interval. Hence, it is an overapproximation of the interval.

□

Theorem 4

If δ is a resolvent of $\Psi_{tr^\lambda(\Pi)}[1, m]$ and $step^\lambda(\delta) = [i, j]$ where $0 \leq i \leq j \leq m$ then δ is a resolvent of $\Psi_{tr^\lambda(\Pi)}[i, j]$ (by lemma 3). For any t $step^\lambda(\delta\langle t \rangle) = [i + t, j + t]$ which means $step^\lambda(\delta\langle t \rangle)$ is a resolvent of $\Psi_{tr^\lambda(\Pi)}[i + t, j + t]$ (by lemma 2). Consequently, for any t where $[i + t, j + t] \subseteq [1, n]$ then $\delta\langle t \rangle$ is entailed by $\Psi_{tr^\lambda(\Pi)}[1, n]$ due to theorem 3.

□

Proposition 3

Let $\mathcal{P}'$ be the set of atoms occurring in some temporal logic program Π that reference the past. For any rule $r \in \Pi^n \cup \Pi^c$ where $B(r) \cap \mathcal{P}' \neq \emptyset$ it holds that $B(r^*) \cap \mathcal{P}' = \emptyset$ since any occurrence is substituted by the corresponding p^* . For any rule $r \in \Pi^n \cup \Pi^c$ where $B(r) \cap \mathcal{P}' = \emptyset$ it holds that $B(r^*) \cap \mathcal{P}' = \emptyset$ since the translation does not change the rule. Hence, for any rule $r \in \Pi^n \cup \Pi^c$ it holds that $B(r^*) \cap \mathcal{P}' = \emptyset$. Which means that $tr^*(\Pi)$ is in PNF.

□

Lemma 4

For any program Π the truth value of p^* and p' always coincide in the resulting program $tr^*(\Pi)$ where $p^* \in \mathcal{P}^*$ are the atoms introduced by the tr^* translation and $p' \in \mathcal{P}'$ are the atoms occurring in Π referencing the past.

Lemma 4

We label the rules added by the tr^* translation as follows:

$$\{p^*\} \leftarrow \quad (A1)$$

$$\perp \leftarrow p', \neg p^* \quad (A2)$$

$$\perp \leftarrow \neg p', p^* \quad (A3)$$

- if p' is True then p^* must also be True to not violate rule A2
- if p' is False then p^* must also be False to not violate rule A3
- if p^* is True then p' must also be True to not violate rule A3
- if p^* is False then p' must also be False to not violate rule A2

We can then conclude that the truth value of p^* and p' always coincide in the resulting program $tr^*(\Pi)$.

□

Proposition 4

We label the rules added by the tr^* translation as follows:

$$\{p^*\} \leftarrow \quad (A4)$$

$$\perp \leftarrow p', \neg p^* \quad (A5)$$

$$\perp \leftarrow \neg p', p^* \quad (A6)$$

Let $p^* \in \mathcal{P}^*$ be the atoms added by the tr^* translation and $p' \in \mathcal{P}'$ be the set of atom occurring in Π that reference the past.

Case 1: Let $\mathcal{P}' \cap A(\Pi) = \emptyset$. Since $tr^*(\Pi) = \Pi$ then $\mathcal{T}_1 = \mathcal{T}_2$ and thus have the same solutions.

Case 2: Let $\mathcal{P}' \cap A(\Pi) \neq \emptyset$. For any solution $\mathcal{S}_1$ of $\mathcal{T}_1$ there is a solution $\mathcal{S}_2$ of $\mathcal{T}_2$ where $\mathcal{S}_2 = \mathcal{S}_1 \cup \{p^* | p' \in \mathcal{S}_1\}$. Since p' and p^* always have the same truth value (lemma 4), the

evaluation of the nogoods induced by $tr^*(\Pi)$ where p' was substituted by p^* will stay the same regardless of the assignment. Also, none of the nogoods induced by the extra rules A5 and A6 will be satisfied since p' and $\neg p^*$ always have different truth values. We can also ignore rule A4 since it does not induce any nogoods. Hence, $\mathcal{S}_2$ is a stable model of $tr^*(\Pi)$ and since it is consistent with I and F it is also a solution to $\mathcal{T}_2$.

On the other hand, For any solution $\mathcal{S}_2$ of $\mathcal{T}_2$ there is a solution $\mathcal{S}_1$ of $\mathcal{T}_1$ where $\mathcal{S}_1 = \mathcal{S}_2 \setminus \{p^* | p' \in \mathcal{S}_2\}$. Since p' and p^* always have the same truth value (lemma 4), the evaluation of the nogoods induced by Π will stay the same regardless of the assignment. Hence, $\mathcal{S}_1$ is a stable model of $\mathcal{T}_1$.

□

Lemma 5

For a resolvent δ of $\Psi_{tr^\lambda(tr^*(\Pi))}[1, m]$ it holds that $simp(\delta)$ is entailed by $\Psi_\Pi[1, m]$ for $\lambda - normal$ solutions.

Lemma 5

Let δ^λ be a resolvent of $\Psi_{tr^\lambda(tr^*(\Pi))}[1, m]$. In $\lambda - normal$ solutions the λ atoms in the nogoods are always true and thus have no effect in their satisfaction. Hence, the nogood $\delta^* = \delta^\lambda \setminus \lambda[1, m]$ is entailed by $\Psi_{tr^*(\Pi)}[1, m]$ since the nogoods in $\Psi_{tr^*(\Pi)}[1, m]$ are the nogoods in $\Psi_{tr^\lambda(tr^*(\Pi))}[1, m]$ without λs .

Let δ^* be a resolvent of $\Psi_{tr^*(\Pi)}[1, m]$ with its corresponding resolution proof $\mathcal{T}$. Let C be the constraints added by the tr^* translation. Observe that a nogood of $tr^*(\Pi) \setminus C$ can be transformed into a nogood of Π simply by substituting all $p^*[i] \in \delta^*$ by their corresponding atom $p[i - 1]$ from $\mathcal{P}$ where $p^* \in \mathcal{P}^*$ are the atoms introduced by the translation and $\mathcal{P}$ is the set of atoms occurring in Π .

For any nogood in $\mathcal{T}$ containing atoms $p^*[i]$, we can substitute them by the corresponding atom $p[i - 1]$ without changing the semantics of the nogoods since they always have the same truth value (by lemma 4).

Next, recall that the constraints added by the tr^* translation have the form

$$\begin{aligned} & \perp \leftarrow p[i - 1], \neg p^*[i] \\ & \text{or} \\ & \perp \leftarrow \neg p[i - 1], p^*[i] \end{aligned}$$

for some integer i . If we substitute p^* by the corresponding atom we get the constraints

$$\begin{aligned} & \perp \leftarrow p[i - 1], \neg p[i - 1] \\ & \text{or} \\ & \perp \leftarrow \neg p[i - 1], p[i - 1] \end{aligned}$$

It is easy to see that any nogood that resolves with the nogoods induced by these constraints would result in the same nogood. Hence, we can remove the nogoods induced by C from $\mathcal{T}$ without affecting its result. Note that the choice rules introduced by the translation do not induce nogoods. This means that any nogood left in $\mathcal{T}$ is either in or entailed by $\Psi_\Pi[1, m]$. Hence, the result δ of the resolution proof $\mathcal{T}$ is entailed by $\Psi_\Pi[1, m]$.

It is clear that $\delta = \{\mathbf{V}p | \mathbf{V}p \in \delta^\lambda, p \in \mathcal{P}\} \cup \{\mathbf{V}p - 1 | \mathbf{V}p^* \in \delta^\lambda, p^* \in \mathcal{P}^*\}$. In words, δ is the result of substituting any atom in $\mathcal{P}^*$ with the corresponding atom in $\mathcal{P}$ and ignoring any λ atoms. Hence, $\delta = \text{simp}(\delta^\lambda)$. Consequently, $\text{simp}(\delta^\lambda)$ is a resolvent of $\Psi_\Pi[1, m]$. $\square$

Theorem 5

Since δ is entailed by $\Psi_{tr^\lambda(tr^*(\Pi))}[1, m]$ then $\text{simp}(\delta)$ is also entailed by $\Psi_\Pi[1, m]$ by lemma 5. By theorem 4 $\text{simp}(\delta)\langle t \rangle$ is entailed by $\Psi_\Pi[1, n]$ for any t where $\text{step}^\lambda(\text{simp}(\delta)\langle t \rangle) \subseteq [1, n]$.

$\square$

Theorem 6

The solution to the temporal logic problem $\langle \Pi, I, F \rangle$ is a stable model of $\text{gen}(\Pi, n)$ consistent with $I[0]$ and $F[0]$. We can split $\text{gen}(\Pi, n)$ as follows: Let $C = \{\{p\} \leftarrow | p \in \mathcal{P}\}$

$$C[0] \cup \Pi[1] \cup \dots \cup \Pi[n]$$

where $\mathcal{P}$ is the set of atoms occurring in Π .

From the Splitting Set Theorem (Lifschitz and Turner 1994) it follows that we can build every stable model X for $\text{gen}(\Pi, n)$ as follows:

X_0 is a stable model of $C[0]$

X_1 is a stable model of $\Pi[1] \cup X_0$

...

X_n is a stable model of $\Pi[n] \cup X_{n-1}$

where X_n is a stable model of $\text{gen}(\Pi, n)$.

It is easy to see that every $X_{i-1} \subseteq X_i$ where $1 \leq i \leq n$. Let $s_i = X_i \cap \mathcal{P}[i]$ with $0 \leq i \leq n$, then program $\Pi[i] \cup X_{i-1}$ can be rewritten as $\Pi[i] \cup s_{i-1} \cup \dots \cup s_0$.

M is a stable model of $\Pi[i] \cup s_{i-1} \cup \dots \cup s_0$ iff M has the form $M_i \cup s_{i-2} \cup \dots \cup s_0$ for some stable model M_i of $\text{trans}(\Pi)[i]$ such that $s_{i-1} = M_i \cap \mathcal{P}[i-1]$. This follows from the fact that $\text{trans}(\Pi)[i] = \Pi[i] \cup C[i-1]$. Following the Splitting Set Theorem, we can build a stable model for $\text{trans}(\Pi)[i]$ by first getting a model S for $C[i-1]$ and then a model for $\Pi[i] \cup S$. Since C is comprised of choice rules for all atoms, then the assignment formed from s_{i-1} is a stable model of $C[i-1]$. Thus, a stable model of $\Pi[i] \cup s_{i-1}$ is a stable model of $\text{trans}(\Pi)[i]$.

This also means that $s_i = M_i \cap \mathcal{P}[i]$ is a state in $G(\Pi)$ and that $\langle s_i, s_{i-1} \rangle$ is an edge.

Consequently, we can say that the states $s_0, \dots, s_n$ form a path in the graph $G(\Pi)$. Finally, for any stable model of $\text{gen}(\Pi, n)$ consistent with $I[0]$ and $F[n]$, then the states $s_0, \dots, s_n$ form a path in $G(\Pi)$ and $I[0]$ and $F[n]$ are consistent with s_0 and s_n respectively.

$\square$

Proposition 5

For simplicity, we consider the case where $\Psi_\Pi[i, j]$ has no body atoms. In the general case, we can convert $\Psi_\Pi[i, j]$ to a set of nogoods with no body atoms that has the same solutions.

By theorem 6 when I and F are empty, the solutions to $\Psi_\Pi[i, j]$ correspond to paths of length $j - i + 1$ in $G(\Pi)$. This means that no path of this length violates a nogood in $\Psi_\Pi[i, j]$.

Since δ is entailed by $\Psi_\Pi[i, j]$, then no path of length $j - i + 1$ in $G(\Pi)$ violates δ .

$\square$

Theorem 7

We prove the case where δ consist of normal atoms. The proof for the general case follows the same lines. Let Π be defined over some set of atoms $\mathcal{P}$. Given that δ is a resolvent of $\Psi_\Pi[i, j]$, its atoms must belong to some smallest set $\mathcal{P}[k, l]$ such that $0 \leq k \leq l$ where $i \leq k + 1$ and $l \leq j$, just like in Proposition 5. Then, the integers t such that $\text{step}(\delta\langle t \rangle) \subseteq [0, n]$ are exactly the t 's such that $-k \leq t \leq n - l$. Hence, to prove this proposition we just have to prove that for every t such that $-k \leq t \leq n - l$ the set of nogoods $\Psi_\Pi[1, n]$ entails $\delta\langle t \rangle$.

Since δ is a resolvent of $\Psi_\Pi[i, j]$, it holds that $\Psi_\Pi[i, j]$ entails δ , and Proposition 5 implies that every path $\langle s_{i-1}, \dots, s_j \rangle$ in $G(\Pi)$ does not violate δ . Since Π is cyclic, the latter implies that every path $\langle s_k, \dots, s_l \rangle$ in $G(\Pi)$ does not violate δ . In turn, this implies that for every path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ and every t such that $-k \leq t \leq n - l$, the subpath $\langle s_{k+t}, \dots, s_{l+t} \rangle$ does not violate the shifted nogood $\delta\langle t \rangle$. Given this and the fact that all literals of $\delta\langle t \rangle$ occur in the subpath $\langle s_{k+t}, \dots, s_{l+t} \rangle$, it follows that for every t such that $-k \leq t \leq n - l$, the paths $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ do not violate $\delta\langle t \rangle$. Then, by Theorem 6, we can conclude that for every t such that $-k \leq t \leq n - l$, the solutions to $\Psi_\Pi[1, n]$ do not violate $\delta\langle t \rangle$, and therefore $\Psi_\Pi[1, n]$ entails $\delta\langle t \rangle$. $\square$

Theorem 8

We prove the case where δ consist of normal atoms. The proof for the general case follows the same lines. We say that a path in $G(\Pi)$ is an initial path if the first state of the path is an initial state of $\langle \Pi, I, F \rangle$. Let Π be defined over some set of atoms $\mathcal{P}$. Given that δ is a resolvent of $\Psi_\Pi[i, j]$, its atoms must belong to some smallest set $\mathcal{P}[k, l]$ such that $0 \leq k \leq l$ where $i \leq k + 1$ and $l \leq j$, just like in Proposition 5. Then, the integers t such that $\text{step}(\delta\langle t \rangle) \subseteq [0, n]$ are exactly the t 's such that $-k \leq t \leq n - l$. Hence, to prove this proposition we just have to prove that for every t such that $-k \leq t \leq n - l$ the nogood $\delta\langle t \rangle$ is not violated by any solution A for $\Psi_\Pi[1, n]$ that is consistent with $I[0]$.

Since δ is a resolvent of $\Psi_\Pi[i, j]$, it holds that $\Psi_\Pi[i, j]$ entails δ , and Proposition 5 implies that every path $\langle s_{i-1}, \dots, s_j \rangle$ in $G(\Pi)$ does not violate δ . Since $\langle \Pi, I, F \rangle$ is cyclic, this implies that every subpath $\langle s_k, \dots, s_l \rangle$ in $G(\Pi)$ of some initial path in $G(\Pi)$ does not violate δ . In turn, this implies that for every initial path $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ and every t such that $-k \leq t \leq n - l$, the subpath $\langle s_{k+t}, \dots, s_{l+t} \rangle$ does not violate $\delta\langle t \rangle$. Given this and the fact that all literals of $\delta\langle t \rangle$ occur in the subpath $\langle s_{k+t}, \dots, s_{l+t} \rangle$, it follows that for every t such that $-k \leq t \leq n - l$, the initial paths $\langle s_0, \dots, s_n \rangle$ in $G(\Pi)$ do not violate $\delta\langle t \rangle$. Then, by Theorem 6, we can conclude that for every t such that $-k \leq t \leq n - l$, the solutions to $\Psi_\Pi[1, n]$ that are consistent with $I[0]$ do not violate $\delta\langle t \rangle$. $\square$

Appendix B Additional results

The following tables show the results of the experiments using the translations from Section 5. The experiments of (Romero et al. 2022) had a bug using those translations in the multi-shot case. Now, that bug is fixed and in those benchmarks the learning approach is no longer worse than the baseline, but it is also not better.

		baseline	500	1000	1500
blocks	(300)	0.5 (0)	0.1 (0)	0.1 (0)	0.1 (0)
depots	(270)	146.4 (30)	138.2 (29)	126.0 (25)	128.3 (30)
driverlog	(135)	14.1 (1)	12.5 (1)	12.3 (1)	10.7 (1)
elevator	(300)	3.0 (0)	3.7 (0)	3.8 (0)	4.3 (0)
grid	(30)	11.4 (0)	5.2 (0)	5.3 (0)	5.2 (0)
gripper	(255)	381.0 (96)	368.5 (91)	359.0 (90)	370.8 (88)
logistics	(225)	0.5 (0)	0.9 (0)	0.9 (0)	0.9 (0)
mystery	(126)	57.0 (3)	58.5 (3)	50.9 (3)	46.8 (2)
Total	(1663)	89.7 (130)	86.4 (124)	82.4 (119)	84.2 (121)

Table B 1. Single shot solving of PDDL benchmarks using translations.

		baseline	500	1000	1500
HanoiTower	(20)	160.6 (2)	97.7 (0)	101.0 (0)	118.2 (1)
Labyrinth	(20)	247.3 (3)	355.7 (4)	355.7 (4)	356.1 (4)
Nomistery	(20)	585.3 (12)	575.6 (12)	556.2 (12)	502.0 (10)
Ricochet Robots	(20)	465.3 (9)	464.7 (9)	464.8 (8)	464.7 (8)
Sokoban	(20)	458.8 (9)	441.5 (9)	458.8 (8)	453.0 (8)
Visit-all	(20)	559.0 (12)	556.5 (12)	560.8 (12)	556.4 (12)
Total	(120)	412.7 (47)	415.3 (46)	416.2 (44)	408.4 (43)

Table B 2. Single shot solving of ASP benchmarks using translations.

		baseline	500	1000	1500
blocks	(20)	1.3 (0)	0.7 (0)	0.7 (0)	0.7 (0)
depots	(18)	148.6 (2)	257.0 (3)	189.4 (3)	221.7 (3)
driverlog	(9)	108.9 (1)	102.0 (1)	104.9 (1)	108.5 (1)
elevator	(20)	280.3 (5)	285.7 (5)	295.0 (5)	305.4 (5)
freecell	(16)	900.0 (16)	900.0 (16)	900.0 (16)	900.0 (16)
grid	(2)	5.2 (0)	4.1 (0)	4.2 (0)	4.3 (0)
gripper	(17)	848.6 (16)	847.5 (16)	849.1 (16)	847.9 (16)
logistics	(20)	225.2 (5)	225.3 (5)	225.3 (5)	225.3 (5)
mystery	(14)	321.8 (5)	321.8 (5)	321.9 (5)	321.9 (5)
Total	(136)	346.6 (50)	361.0 (51)	353.8 (51)	359.7 (51)

Table B 3. Multi shot solving of PDDL benchmarks using translations.

		baseline	500	1000	1500
HanoiTower	(20)	554.1 (10)	601.4 (11)	593.7 (10)	646.7 (11)
Labyrinth	(20)	647.7 (14)	647.8 (14)	647.8 (14)	647.9 (14)
Nomistery	(20)	64.2 (1)	77.0 (1)	81.0 (1)	69.3 (1)
Ricochet Robots	(20)	527.3 (11)	518.1 (11)	519.3 (11)	521.3 (11)
Sokoban	(20)	721.5 (16)	722.6 (16)	722.3 (16)	722.0 (16)
Visit-all	(20)	677.5 (13)	704.0 (13)	774.6 (15)	801.6 (16)
Total	(120)	532.1 (65)	545.2 (66)	556.5 (67)	568.1 (69)

Table B 4. Multi shot solving of ASP benchmarks using translations.