

A Practical Approach to Discretised PDDL+ Problems by Translation to Numeric Planning

Francesco Percassi

*School of Computing and Engineering
University of Huddersfield, UK*

F.PERCASSI@HUD.AC.UK

Enrico Scala

*Dipartimento di Ingegneria dell'Informazione
Università degli Studi di Brescia, Italy*

ENRICO.SCALA@UNIBS.IT

Mauro Vallati

*School of Computing and Engineering
University of Huddersfield, UK*

M.VALLATI@HUD.AC.UK

Abstract

PDDL+ models are advanced models of hybrid systems and the resulting problems are notoriously difficult for planning engines to cope with. An additional limiting factor for the exploitation of PDDL+ approaches in real-world applications is the restricted number of domain-independent planning engines that can reason upon those models.

With the aim of deepening the understanding of PDDL+ models, in this work, we study a novel mapping between a time discretisation of PDDL+ and numeric planning as for PDDL2.1 (level 2). The proposed mapping not only clarifies the relationship between these two formalisms but also enables the use of a wider pool of engines, thus fostering the use of hybrid planning in real-world applications. Our experimental analysis shows the usefulness of the proposed translation and demonstrates the potential of the approach for improving the solvability of complex PDDL+ instances.

1. Introduction

The availability of domain-independent planning engines is fostering the use of automated planning in a wide range of applications. This is despite the complexity issues inherent in plan generation, which are exacerbated by the separation of planning logic from domain knowledge (McCluskey & Porteous, 1997). A major advantage of the separation of planning logic from domain knowledge lies in the fact that, given a standard language to be used for input and output, the two components can be interchanged in a modular way, without affecting the other component, and with no negative repercussions on the overall application framework where the planning system is usually embedded (McCluskey, Vaquero, & Vallati, 2017; Vallati, Chrapa, McCluskey, & Hutter, 2021).

This modular approach promotes the use of reformulation techniques, which can automatically re-formulate or re-represent the domain knowledge and/or the problem instance in order to increase the effectiveness of the planning logic component and increase the scope of problems solved. The general idea is to develop reformulation techniques that are agnostic about the domain knowledge and the planning logic and use them to form a wrapper around the planning engine, improving its performance for the domain in which it is applied. The transformation is then reversed after a solution has been found, such

that the solution is rephrased in the original formulation language. A well-known class of reformulation approaches aims at translating a model from the original input language to a different one. The idea is usually to remove the use of some poorly supported features of the language (Helmert, 2009; Percassi & Gerevini, 2019; Bonassi, Gerevini, Percassi, & Scala, 2021) or to re-represent the problem in a less expressive language. The latter strategy has the advantage of increasing the number of planning engines that are able to reason upon the planning problem, and leverage existing robust technologies devised for solving more restricted cases. Well-known examples of this approach include the translation of conformant planning problems into classical problems (Taig & Brafman, 2013; Grastien & Scala, 2020), the re-representation of uncertainty in conformant planning problems (Palacios & Geffner, 2009), and the translation of complex temporal aspects in PDDL2.1 (Cooper, Maris, & Régnier, 2010).

In this paper, we introduce a reformulation approach for translating discretised PDDL+ planning problem instances into PDDL2.1 instances. PDDL+ models are advanced models of hybrid systems and the resulting problems are notoriously difficult to cope with. Further, a limited number of planning engines are able to parse PDDL+ models, therefore translating into PDDL2.1 can significantly extend the pool of planning engines that can be used to solve a given instance.

More precisely we study two translations. The first translation leads to a numeric planning problem that is exponentially larger than the input PDDL+ but preserves the number of discrete transitions. The second one keeps the resulting formulation polynomial but requires also a polynomial number of additional transitions to generate a solution. We start off by revisiting the formalisation provided by Shin and Davis (2005), which lets us formalise the problem without the need to depend on the hybrid automaton interpretation proposed by Fox and Long (2003), and also crisply formalises the connection between the continuous-time representation, and its discretisation. We formally present the two translations and show how these can be extended to encode a cascade of events emulated by actions. Further, we introduce two approaches for optimising the operability of the translated models; this makes the resulting PDDL2.1 models more effective. Indeed, our optimisations let us reduce the number of necessary checks along with a plan whilst preserving soundness and completeness with a pre-processing phase that looks at the structure of our problem instance and applies the optimisations only when it is possible. We validate the resulting formulations against a set of challenging benchmark domains, including real-world applications, and well-known planning engines, and we assess the impact of the introduced optimisations. Our results indicate that the proposed translations can unlock the use of PDDL2.1 planning engines for tackling hybrid PDDL+ problems, with the clear advantage of expanding the number of approaches that can be used to solve a problem.

This paper significantly extends our conference paper on the topic (Percassi, Scala, & Vallati, 2021). This extended version advances our previous work along several dimensions. Firstly, we revise the formalisation of PDDL+ and its discretisation that are not fully spelt out in the conference paper. We enrich our discussion with examples and more formal definitions. This reformulation does not change the meaning of our previous formalisation but makes things more precise and less ambiguous. Secondly, we provide full proofs for all the translations that we have presented and also enrich the discussion with a concrete full example that is an extension of the well-known car domain by Fox and Long (2006).

To these proofs, we also add new theorems on the size of our translated output. Thirdly, we formalise and study theoretically two optimisations that we apply to our translation schemata. Such optimisations are meant to reduce the length of the plans obtained on the translated output. Last but not least, we substantially extend our experimental analysis to evaluate our contribution, pros and cons, in greater detail.

This paper is organised as follows. Section 2 formalises the problem of PDDL+. The proposed translations are then presented in detail in Section 3. Section 4 focuses on different ways for optimising the models generated by the introduced translations. An extensive experimental analysis is provided in Section 5. Related works are discussed in Section 6. Finally, Section 7 frames the scope of the proposed approach, highlighting its advantages and discussing some limitations of the general discretisation-based approach, and conclusions are given in Section 8.

2. Problem Formalisation

In this section we formalise the problem of PDDL+ (Fox & Long, 2006) without durative actions, thus focusing on the core features of this formalism, and the problem of numeric planning as the one that can be specified in PDDL2.1, Level 2 (Fox & Long, 2003), hereinafter simply referred to as PDDL2.1. We first describe the syntax of our problems and then detail the semantics. Our discussion follows the formalisation and terminology provided by Shin and Davis (2005) in a way that is instrumental for our work.

Let F be a set of Boolean variables, and X be a set of numeric variables taking values from $\mathbb{R}_? = \mathbb{R} \cup \{f, ?\}$; $?$ represents an undefined numeric value. A PDDL+ state s is a full assignment to all variables in F and in X . Given a state s and a variable $v \in F \cup X$, we make use of $s[v]$ to indicate the value assumed by v in s . This value can either be a numeric value or a Boolean value depending on whether v is in X or in F , respectively. A numeric expression is defined inductively as follows: $x \in X$, and $k \in \mathbb{Q}$ are numeric expressions; let ξ and ξ' be two numeric expressions, ξ *op* ξ' with *op* $\in \{+, -, \cdot, /, \%$ $\}$, g is a numeric expression. For instance, $x + 5 + y$ is a numeric expression while $x + +5$ is not. Let ξ be a numeric expression, we denote as $s[\xi]$ the evaluation of ξ in s . A division between two numeric expressions ξ and ξ' , i.e., ξ/ξ' , is undefined in s if $s[\xi'] = 0$, i.e., division by zero is undefined. Moreover, let x be a variable involved in an expression ξ . If x is undefined in state s , then ξ is undefined too.

Boolean and numeric assignments both have the form $hf := bi$. For Boolean assignments, f is a Boolean variable and $b \in \mathbb{B} = \{f, ?\}$. For numeric assignments, $hf \text{ asgn, inc, dec } g, x, \xi$, where *fasgn, inc, dec* are the contractions of the keywords *assign, increase* and *decrease*, respectively, x is a numeric variable and ξ is a numeric expression, respectively. A Boolean condition has the form $hf = bi$, where f is a Boolean variable and $b \in \mathbb{B}$. A numeric condition has the form $h\xi \bowtie 0$ where ξ is a numeric expression and $\bowtie \in \{f, <, =, >\}$, g . We detail our problems using formulas expressed in Negation Normal Form (NNF). The terms of our formulas contain both propositional and numeric conditions.

Definition 1 (PDDL+ Problem). *A PDDL+ planning problem Π is a tuple $\langle F, X, I, G, A, E, P \rangle$ where:*

- *F and X are sets of Boolean and numeric variables, respectively;*

- I is the description of the initial state expressed as a full assignment to all variables in X and F ;
- G is the description of the goal expressed as a formula;
- A and E are the sets of actions and events, respectively. Actions and events are pairs hp, ei where p is a formula and e is a set of conditional effects of the form $c \triangleright e$ where (i) c is a formula and (ii) e is a set of Boolean or numeric assignments;
- P is a set of processes. A process is a pair hp, e^0i where p is a formula and e^0 is a set of numeric continuous effects expressed as pairs $hx, \xi j$ where $x \subseteq X$ and ξ is a numeric expression.¹

Let $a = hp, ei$ be an action or an event or a process, we use $pre(a)$ to denote the precondition p of a , and $eff(a)$ the effect e of a . In what follows we generally use the symbols a , ρ and ε for denoting an action, a process, and an event, respectively. Given a PDDL+ problem Π we denote the whole state space as $states(\Pi) = \mathcal{B}^{X \cup F} \times \mathcal{R}^{X \cup J} \cup \{undefined\}$, where $undefined$ denotes a special state that represents the notion of inconsistency within our semantic (this aspect will be dealt with in detail later). Note that, although X takes values in \mathbb{R} , it is not possible to use real values in PDDL+ in the problem specification.

For ease of exposition and following PDDL common practice, in our examples and encodings we will, when convenient, represent the state and the effects using a set-theoretic representation. Boolean conditions and assignments having the extended form $hf = >i$ ($hf := >i$) and $hf = ?i$ ($hf := ?i$) are shortened, by abuse of notation, to f and $:f$, and a conditional effect where the left-hand side is always satisfied i.e., $>\triangleright e$, is rewritten as e . For example, formula $ha = >i \wedge hb = ?i$ is contracted into $a \wedge :b$, and an assignment having the form $fha := >i, hb := ?ig$ is contracted into $fa, :bg$. In the set-theoretic representation, a state is expressed as a pair encompassing a set of facts and a set of assignments to numeric variables: if a fact is not mentioned in the state, the Boolean variable associated with that fact is false (closed world assumption); similarly, if a numeric variable is not given any assignment in the state, such a variable is undefined. For example, a state $s = fha = >i, hb = ?ig$, where a and b are Boolean variables, can be expressed more succinctly by $s = fa, g$.

A plan for a PDDL+ problem is an ordered set of timed actions plus a time envelope, organised formally as the following.

Definition 2 (PDDL+ Plan). A PDDL+ plan π_t for a PDDL+ problem Π is a pair $\langle \pi, \langle t_s, t_e \rangle \rangle$ where $\pi = \langle \langle ha_0, t_0 \rangle, \langle ha_1, t_1 \rangle, \dots, \langle ha_{n-1}, t_{n-1} \rangle \rangle$ is a finite sequence of pairs such that for each $i \in [0..n-1]$ we have that $a_i \in A$ and $t_s \leq t_i \leq t_e$; $\langle t_s, t_e \rangle$, with $t_s, t_e \in \mathbb{Q}$ and $t_s \leq t_e$, is the envelope within which π is performed. π_t is a well-formed plan if $\exists i, j \in [0..n-1]$ and $i < j$, then $t_i \leq t_j$ holds.

Hereinafter, we consider just well-formed plans. Let $\pi_t = \langle \pi, \langle t_s, t_e \rangle \rangle$ be a PDDL+ plan and let $\langle ha_i, t_i \rangle$ and $\langle ha_j, t_j \rangle$ be two pairs appearing in π . We write $\langle ha_i, t_i \rangle \prec \langle ha_j, t_j \rangle$ ($a_i \prec a_j$)

1. We will see that under continuous time interpretation, \dot{x} denotes the net derivative of x .
 2. $[1..n]$ is the integer interval $\{1, \dots, n\}$.

iff $ha_i, t_i (a_i)$ appears strictly before $ha_j, t_j (a_j)$ in π . Furthermore, given a sequence of objects A we denote its set representation as $set(A)$, e.g., $set(\pi)$ is the set of all time-stamped actions in π .

Having formally defined the PDDL+ problem, we now see how a PDDL2.1 one is defined.

Definition 3 (PDDL2.1 Problem). *A PDDL2.1 planning problem Π is a tuple $\langle F, X, I, G, A, c \rangle$ where all elements are as for PDDL+, yet there are neither processes nor events, and c is a function that associates to each action a rational non-negative cost.*

We clarify that our definition of PDDL2.1 slightly differs from that provided by Fox and Long (2003) in which action costs are not included.

For a PDDL2.1 problem Π , the state space is defined as $states(\Pi) = \mathcal{P}(\mathcal{B}^{F \cup X}) \setminus \{g \mid g \text{ undefined}\}$.

Definition 4 (PDDL2.1 Plan). *A PDDL2.1 plan for a PDDL2.1 problem Π is simply a finite sequence of actions from A .*

Again, let π be a PDDL2.1 plan and let a_i and a_j be two actions appearing in π . We write $a_i \prec a_j$ iff a_i appears strictly before a_j in π .

Before going into the formal details of the semantics of PDDL+ and PDDL2.1, let us briefly explain intuitively what these problems actually entail. Roughly speaking, a PDDL+ problem consists in finding a number of actions along a potentially infinite timeline, whilst conforming to a number of processes and events that may change the state of the world in a continuous or instantaneous manner as time goes by. A solution for a PDDL+ problem is such iff all actions are applicable when they are scheduled and the goal is met at time t_e . In technical terms, PDDL+ prescribes events and processes to be interpreted as *must* discrete and continuous transitions along a potentially infinite timeline, while actions are *may* transitions. A PDDL2.1 problem is a variant where there is no explicit management of time and exogenous changes, and we only seek a sequence of actions that starts from some initial state and yields a state satisfying the goal.

We start with the simpler semantics of PDDL2.1 and then delve into the details of PDDL+ by stressing, in particular, its temporal aspect.

The satisfiability of a formula follows the usual semantics of propositional logic extended with numeric conditions. In particular, let $h\xi \bowtie 0i$ be a numeric condition, $s \models h\xi \bowtie 0i$ iff $s[\xi] \bowtie 0$ evaluates to true. A formula containing a numeric expression that is undefined in the state s is treated as unsatisfiable in s by default reasoning. For instance, if we have $\alpha = hx \leq 0i \wedge hx \leq 10i$ with x undefined in s , then α is not satisfied in s ; similarly, let $\beta = hx \leq 0i \wedge hy = 1i$ we have that $s \models \beta$ iff $s[y] = 1$.

The application of a in s yields the state $s^0 = \Gamma(s, a)$ in a way that, for each $v \in F \cup X$, the following holds:

$$\Gamma(s, a) = \begin{cases} s^0[v] = b & \text{if } \exists \alpha \triangleright f\dots, hv := b \text{ and } s \models \alpha, \text{ where } b \in \mathbb{B}; \\ s^0[v] = s[\xi] & \text{if } \exists \alpha \triangleright f\dots, hasgn, v, \xi \text{ and } s \models \alpha; \\ s^0[v] = s[v] + s[\xi] & \text{if } \exists \alpha \triangleright f\dots, hinc, v, \xi \text{ and } s \models \alpha; \\ s^0[v] = s[v] - s[\xi] & \text{if } \exists \alpha \triangleright f\dots, hdec, v, \xi \text{ and } s \models \alpha; \\ \text{undefined} & \text{if there exists at least a pair of conflicting assignments in } eff(a); \\ s^0[v] = s[v] & \text{otherwise.} \end{cases}$$

Case (i) of $\Gamma(s, a)$ deals with the propositional assignment, (ii) the simple numeric assignment, (iii) and (iv) the additive numeric assignments. Case (v) handles the possibility that two contradicting effects exist, either Boolean or numeric. For a given state s and an action a , two Boolean effects $\alpha \triangleright f\dots, hf := >ig, \beta \triangleright f\dots, hf := ?ig \ 2 \ \text{eff}(a)$, where α and β are two formulae and $f \ 2 \ F$, are conflicting if $s \not\models \alpha \wedge \beta$, and two numeric effects $\alpha \triangleright f\dots, hop, x, \xi ig, \beta \triangleright f\dots, hop^l, x, \xi^l ig \ 2 \ \text{eff}(a)$ with $x \ 2 \ X$ and $op, op^l \ 2 \ \text{fasgn}, inc, decg$, are conflicting if $s \not\models \alpha \wedge \beta$; in other words, two numeric effects are conflicting if they contextually affect the same variable in different ways. Finally, case (vi) is the frame-axiom, so the value of v persists if no effect of a affects v .

Let $ha_0, \dots, a_{n-1}i$ be a sequence of actions, we say that it is applicable in s iff $s \models \text{pre}(a_0)$, $\Gamma(s, a_0) \models \text{pre}(a_1)$, \dots , $\Gamma(\dots \Gamma(s, a_0), a_{n-2}) \models \text{pre}(a_{n-1})$. We use $\gamma(s, \cdot)$ for the state resulting by applying either an action, i.e., $\gamma(s, a)$, or a sequence of (applicable) actions, i.e., $\gamma(s, ha_0, \dots, a_{n-1}i)$, in state s . Since actions and events have the same structure, we use the same notation for denoting the state resulting from the application of an event, i.e., $\gamma(s, \varepsilon)$, or a sequence of (applicable) events, i.e., $\gamma(s, h\varepsilon_0, \dots, \varepsilon_{n-1}i)$.

Definition 5 (Valid and Optimal PDDL2.1 Plan). *Let $\Pi = \langle hF, X, I, G, A, ci \rangle$ be a PDDL2.1 problem, a plan $\pi = ha_0, \dots, a_{n-1}i$ is said to be a valid plan for Π if it is applicable and $\gamma(I, \pi) \models G$. The plan π is said to be optimal if among all valid plans for Π , it is the one that minimises the total cost, $\pi = \underset{\text{valid for } \Pi}{\text{argmin}} \sum_{a \in \text{set}(\pi)} c(a)$.*

The semantics of PDDL+ can be specified through the notion of time points, intervals, and histories over intervals. We use these to characterise the projection of a plan and conclude with the formalisation of the validity of a PDDL+ plan which we interpret both on a continuous and a discrete timeline. We start with the continuous version and then introduce its discretisation. Note that all our translations that we present from Section 3 are all assuming a discrete model.

Definition 6 (Time Point). *A time point T is a pair ht, ni where $t \ 2 \ \mathbb{R}$ and $n \ 2 \ \mathbb{N}$.*

Given a time point $T = ht, ni$, in the rest of the paper we will refer to t as the clock and n as the step of T .

Time points over $\mathbb{R} \times \mathbb{N}$ are ordered lexicographically, i.e., let ht_1, n_1i and ht_2, n_2i be two time points, $ht_1, n_1i < ht_2, n_2i$ iff either $t_1 < t_2$ or $t_1 = t_2$ and $n_1 < n_2$. Let T_1 and T_2 be two time points, a closed (open) time interval $l = [T_1, T_2]$ ((T_1, T_2)) is the non-empty set $l = \{T \mid T_1 \leq T \leq T_2\}$ ($\{T \mid T_1 < T < T_2\}$).

Intuitively, a time point allows us to order situations along time and impose an ordering, by using a natural number n , among two different situations whenever such situations share the same clock. This representation idealises the execution of actions and events to be truly instantaneous as long as an order is imposed among transitions sharing the same clock as done in the work of Shin and Davis (2005). Note that PDDL+ (Fox & Long, 2006) supports such a temporal model for events only; mutexes actions instead are still required to be temporally separated by at least a small quantity of time ϵ (ϵ -separation requirement).³

3. For a deeper analysis on the temporal implications of ϵ - vs non- ϵ -separation requirement in the context of temporal PDDL2.1, the interested reader can look at the paper by Gigante, Micheli, Montanari, and Scala (2022).

Definition 7 (History). A history H for a PDDL+ problem Π over $I = [T_s, T_e]$ maps each time point in I into a situation. A “situation at time point T ” is the tuple $H(T) = \langle H_A(T), H_S(T) \rangle$, where $H_A(T)$ is a sequence of actions executed at time point T and $H_S(T)$ is a state, i.e., the assignment to all variables in X and F at time point T .

An event ε is active (triggered) in a state s iff $s \models \text{pre}(\varepsilon)$. Given a history H over $I = [T_s, T_e]$ and a time point $T \in I$, multiple events can be triggered.

Hereinafter, given a set S , we denote with $\text{seq}(S)$ any possible sequencing of its elements.

Definition 8 (Triggered Events). Let H be a history for a PDDL+ problem Π over I and let $T \in I$. The sequence of events triggered at time point T is defined as $E_{\text{trigg}}(T) = \text{seq}(\{ \varepsilon \in E \mid H_S(T) \models \text{pre}(\varepsilon) \})$.

A time point is said to be significant if something meaningful happens, such as the execution of an action, the triggering of a sequence of events, or a change in the set of active processes.

Definition 9 (Significant Time Point). $T = \langle t, n \rangle$ is a significant time point (hereinafter STP) of a history H for Π over $I = [T_s = \langle t_s, n_s \rangle, T_e = \langle t_e, n_e \rangle]$, iff $T \in I$ and at least one of the following holds:

1. $T = T_s$ or $T = T_e$; the time points associated with the extremes of the time envelope are always STPs;
2. $H_A(T) \neq \langle \rangle$;
3. $E_{\text{trigg}}(T) \neq \langle \rangle$, i.e., there exists at least an event $\varepsilon \in E$ such that $H_S(T) \models \text{pre}(\varepsilon)$;
4. there has been a discrete change just before; formally it has to hold that $n > 0$ and there exists a $T^0 = \langle t, n-1 \rangle$ such that $H_A(T^0) \neq \langle \rangle$ or $E_{\text{trigg}}(T^0) \neq \langle \rangle$;
5. a process has started (stopped) in T ; formally $H_A(T) = \langle \rangle$, $E_{\text{trigg}}(T) \neq \langle \rangle$ and there exists $\rho \in P$ for which $H_S(T) \models \text{pre}(\rho)$ ($H_S(T) \not\models \text{pre}(\rho)$), there exists $T^0 < T$ such that $T^0 \in I$ and for each T^{00} such that $T^0 > T^{00} > T$ then $H_S(T^{00}) \not\models \text{pre}(\rho)$ ($H_S(T^{00}) \models \text{pre}(\rho)$).

We remark that Point 5 of the previous definition, i.e. the condition that determines the starting (stopping) of a process, is to be interpreted as a high-level definition and not an operational one. Specifically, in the case of continuous semantics discussed therein, identifying the exact point where a context switch occurs can be pointless even in simple cases. For example, consider the case in which the precondition of a process requires $hx > 0$ to be activated and x is linearly incremented starting from 0. For handling cases like this, it is necessary to specialise the provided condition by considering the case in which the context switch involves a closed or open condition. Specifically, in the second case, for providing an operational definition it is necessary to exploit the concept of mathematical limit. As this work focuses primarily on discrete semantics, such a level of detail is outside its scope.

Within an interval, we can also identify a number of sub-intervals where nothing meaningful happens. This is the notion of a monotonous interval. Formally:

Definition 10 (Monotonous interval). *A history H of Π over $I = [T_s = ht_s, n_s i, T_e = ht_e, n_e i]$ is monotonous over $I_t = (t_1, t_2) \quad (t_s, t_e)$, where $t_1, t_2 \geq \mathbb{R}$, if for each $t \geq I_t$ and for any given n, ht, ni is not a STP of H .*

Note that, for each ht, ni such that $t \geq (t_1, t_2)$, the set of active processes does not change and the sequence of triggered events is empty. We denote the set of active processes over I_t as the context $C(I_t) = \{ \rho \geq P \mid H_s(ht_1, n_1 i) \not\models pre(\rho) \}$, where n_1 is a sufficiently large natural number beyond which the state is stable, i.e., for each step $n \geq n_1$ then $H_A(ht_1, n i) = hi$ and $E_{trigg}(ht_1, n i) = hi$. A context switch happens when the set of active processes changes.

In order to guarantee that there is only a unique and finite set of significant time points given a history, it is necessary to impose some restrictions on its structure. In what follows we will go over the issues caused by events and context switches, and conclude with a number of assumptions that are needed to have a meaningful and manageable notion of plan projection, which is what we use to infer whether a plan is valid or not. Similar restrictions have been used also in other works on PDDL+ (Fox, Howey, & Long, 2005; Shin & Davis, 2005; Cashmore, Magazzeni, & Zehtabi, 2020).

We start off this discussion with an example showing how the general formulation of PDDL+ can induce non-deterministic spontaneous transitions.

Example 2.1 (Non-deterministic Events). *Let $\varepsilon_1 = hhx = 0i, fhasgn, y, 1i, hasgn, x, 1igi$ and $\varepsilon_2 = hhx = 0i, fhasgn, y, \quad 1i, hasgn, x, 1igi$ be two events and let s be a state such that $s \not\models hx = 0i$. The set of triggered events in s is $\{ \varepsilon_1, \varepsilon_2 \}$, since $s \not\models pre(\varepsilon_1) \wedge pre(\varepsilon_2)$. Such a set can be sequenced into $h\varepsilon_1, \varepsilon_2i$ or $h\varepsilon_2, \varepsilon_1i$ and, depending on the chosen ordering, two different outcomes can be obtained, i.e., $s^0 = \gamma(s, h\varepsilon_1, \varepsilon_2i) = \gamma(\gamma(s, \varepsilon_1), \varepsilon_2) \not\models hy = 1i$ and $s^{00} = \gamma(s, h\varepsilon_1, \varepsilon_2i) = \gamma(\gamma(s, \varepsilon_2), \varepsilon_1) \not\models hy = \quad 1i$. A problem including these events is said to be non-deterministic as different sequences of triggered events can produce different resulting states.*

To avoid this possibility, Fox and Long (2006) provided the definition of an event-deterministic PDDL+ problem. This definition requires that in any state where more than one event is triggered, the outcome resulting from their sequential execution is independent of their ordering.

Definition 11 (Event-deterministic PDDL+ Problem (Fox & Long, 2006)). *A PDDL+ problem Π is said to be event-deterministic if, for each state where two events ε_1 and ε_2 are triggered, then the transition sequences induced by $h\varepsilon_1, \varepsilon_2i$ and $h\varepsilon_2, \varepsilon_1i$ are both valid and reach the same state. In this case, ε_1 and ε_2 are said to commute.*

Determining whether a PDDL+ problem is event-deterministic or not may be expensive, but it is possible to establish a sufficient criterion that ensures that this always holds. For example, two non-mutex events always commute. Indeed, if every pair of events that are triggered in any state commute, then the problem is event-deterministic. In the following, we assume that PDDL+ problems are event-deterministic, which implies the guarantee that any order of evaluation chosen from the triggered events in a state leads to the same state.

Another source of complexity comes with a potentially infinite cascade of events. A cascade of events occurs when a state $H_s(ht, ni)$ triggers at least an event, i.e., $E_{trigg}(ht, ni) =$

$h\varepsilon_1 i$, such that the resulting state obtained by the application of ε , i.e., $H_S(ht, n + 1) = \gamma(H_S(ht, n), h\varepsilon_1 i)$, triggers again at least an event, i.e., $H_S(ht, n + 1) \not\models pre(\varepsilon_2)$.

Hereinafter, given two sequences of elements V and V^0 , we denote as $V \circ V^0$ the sequential merging of V and V^0 . For example, given $V = h1, 2i$ and $V^0 = h2, 3i$, then $V \circ V^0 = hh1, 2i, h2, 3ii$. The \circ operator is associative and the identity element is the empty sequence, so $V \circ hi = V$.

Definition 12 (Cascade of Events). *Let Π be a PDDL+ problem, H be a history for Π over l and let $T = ht, ni \geq l$. The cascade of events triggered in T is defined recursively as a sequence of sequences as follows:*

$$ES(T) = \begin{cases} hi & \text{if } E_{trigg}(T) = hi \\ E_{trigg}(ht, ni) \circ ES(ht, n + 1) & \text{otherwise} \end{cases}$$

To give an intuition of the definition, look at the following small example.

Example 2.2 (Cascade of Events). *Let Π be a PDDL+ problem encompassing three events, $E = f\varepsilon_1, \varepsilon_2, \varepsilon_3g$, and $\varepsilon_1 = hhx = 0i, fhasgn, x, 1igi$, $\varepsilon_2 = hhx = 1i, fhasgn, x, 2igi$ and $\varepsilon_3 = hhx = 2i, fhasgn, x, 3igi$. Let H be a history for Π over l and let $T = ht, ni \geq l$ be a STP such that $H_S(ht, ni) \not\models hx = 0i$. Then, by iterating recursively over the states obtained executing the events, we get that:*

- $H_S(ht, ni) \not\models hx = 0i \not\models pre(\varepsilon_1) \wedge pre(\varepsilon_2) \wedge pre(\varepsilon_3)$, then $E_{trigg}(ht, ni) = h\varepsilon_1 i$;
- $H_S(ht, n + 1) = \gamma(H_S(ht, ni), E_{trigg}(ht, ni) = h\varepsilon_1 i) \not\models hx = 1i \not\models pre(\varepsilon_1) \wedge pre(\varepsilon_2) \wedge pre(\varepsilon_3)$, then $E_{trigg}(ht, n + 1) = h\varepsilon_2 i$;
- $H_S(ht, n + 2) = \gamma(H_S(ht, n + 1), E_{trigg}(ht, n + 1) = h\varepsilon_2 i) \not\models hx = 2i \not\models pre(\varepsilon_1) \wedge pre(\varepsilon_2) \wedge pre(\varepsilon_3)$, then $E_{trigg}(ht, n + 1) = h\varepsilon_3 i$;
- $H_S(ht, n + 3) = \gamma(H_S(ht, n + 2), E_{trigg}(ht, n + 2) = h\varepsilon_3 i) \not\models hx = 3i \not\models pre(\varepsilon_1) \wedge pre(\varepsilon_2) \wedge pre(\varepsilon_3)$, then $E_{trigg}(ht, n + 3) = hi$.

The cascade of events triggered in T is defined as $ES(T) = E_{trigg}(ht, ni) \circ ES(ht, n + 1) = E_{trigg}(ht, ni) \circ E_{trigg}(ht, n + 1) \circ ES(ht, n + 2) = E_{trigg}(ht, ni) \circ E_{trigg}(ht, n + 1) \circ E_{trigg}(ht, n + 2) \circ ES(ht, n + 3) = h\varepsilon_1 i \circ h\varepsilon_2 i \circ h\varepsilon_3 i \circ hi = hh\varepsilon_1 i, h\varepsilon_2 i, h\varepsilon_3 ii$.

As hinted at above, Definition 12 allows the characterisation of an infinite cascade of events. Fox et al. (2005) discusses limitations to be imposed on the structure of events, in addition to those described above concerning the non-determinism, in order to prevent such an infinite cascade of events from actually happening. The rule is that every event has to be *self-deactivating*. An event is self-deactivating if it falsifies its preconditions when it is applied, i.e., for each state $s \not\models pre(\varepsilon)$ then $\gamma(s, \varepsilon) \not\models pre(\varepsilon)$. This is necessary to prevent a triggered event from continuing to trigger indefinitely. It should be noted that there can also be more complex cyclic-triggering situations involving more than a single event. For this reason, it is assumed that an event is triggered at most once in a timestamp (Fox et al., 2005).

Our assumptions construct on the above notions and can be then summarised as follows:

Assumption 1 (Event-determinism). *The PDDL+ problem Π has to be event-deterministic (see Definition 11).*

Assumption 2 (Finite Complexity). *The PDDL+ problem Π has to induce a finite number of spontaneous changes over an interval I_t . This finiteness is imposed both on the number of events that can be triggered, and the number of times the starting and the stopping of a process causes a change of context.*

It is easy to see that the above assumptions lead us to get a unique and finite number of STPs over an interval I_t .

We are now ready to define the projection of a PDDL+ plan. Intuitively, we define the plan projection around a number of rules that specify how history evolves over time. The first (second) rule states that if an action (event) is executed (triggered) in an STP $T = ht, ni$, then there necessarily exists a successor of T , i.e., $T^\theta = ht, n+1i$ having the same clock t and the step increased by one unit, i.e., $n+1$; the successor state associated to T^θ is calculated by simply applying the discrete effects of the actions (events). The third rule is used to enforce how actions of a PDDL+ plan π are projected over a history, preserving their original ordering in case they share the same timestamp in π . Then, the fourth rule is used to enforce how a numeric variable changes continuously over time according to the active processes in those “monotonous” temporal intervals in which “nothing happens” (there is no action/event executed/triggered and there is no process which starts/ends). The following definition formalises this intuition, assuming $\Pi = hF, X, I, G, A, E, Pi$, and $\pi_t = h\pi, ht_s, t_{eii}$ implicit.

Definition 13 (PDDL+ Plan Projection). *Let H be a history for Π over I , let I be an initial state and let π_t be a PDDL+ plan for Π . We say that H is a projection of π_t which starts in I iff H induces a finite sequence of STPs $T_H = hT_0 = ht_s, 0i, \dots, T_m = ht_e, n_mii$ such that H is defined over $I = [T_0, T_m]$ with $H_S(T_0) = I$, $H_A(T_m) = hi$, $E_{trigg}(T_m) = hi$ and, for all $i \in [0..m]$, the following rules hold:*

R1 $E_{trigg}(T_i) \neq hi$ iff $H_S(T_{i+1}) = \gamma(H_S(T_i), E_{trigg}(T_i))$, $H_A(T_i) = hi$, $t_{i+1} = t_i$ and $n_{i+1} = n_i + 1$;

R2 $H_A(T_i) \neq hi$ iff $H_S(T_{i+1}) = \gamma(H_S(T_i), H_A(T_i))$, $E_{trigg}(T_i) = hi$, $t_{i+1} = t_i$ and $n_{i+1} = n_i + 1$;

R3 for each ha_i, t_i, ha_j, t_j in π , with $i < j$ and $t_i = t_j$ there exists T_k, T_z in T_H such that a_i in $H_A(T_k)$ and a_j in $H_A(T_z)$ where $t_k = t_z = t_i$ and $n_k < n_z$;

R4 H is monotonous over $I_t = (t_i, t_{i+1})$, with $t_i < t_{i+1}$ and for each $x \in X$, we have that:

- $H_S(ht, 0i)[x]$ is continuous and differentiable over I_t ;
- $\forall t \in I_t$ it holds that

$$\frac{dH_S(ht, 0i)[x]}{dt} = \sum_{\substack{hx^\theta; i \in \mathbb{Z}; x^\theta = x \\ \mathbb{Z} \subset I_t}} H_S(ht, 0i)[\xi]$$

- $H_S(ht_{i+1}, 0i)[x] = \lim_{t \downarrow t_{i+1}} H_S(ht, 0i)[x]$, and values of unaffected variables persist up to t_{i+1} (frame-axiom).

To better understand Definitions 6–13, the example of Figure 1 shows how a numeric variable, in this case, the velocity of a car, changes continuously over time according to a PDDL+ plan. It is remarkable how the PDDL+ problems can induce dynamics on numeric variables having piece-wise defined dynamics with points of discontinuity, as shown for instance in Figure 1. This is why monotonous intervals must necessarily be defined as open intervals and the numeric variable must be continuous and time-differentiable only within them (R4 of Definition 13).

Finally, we remark that, as R3 is defined, for each $T \geq l$ then $jH_A(T)j = 1$.

Definition 14 (Valid PDDL+ Plan). *Let Π be a PDDL+ problem. Let π_t be a PDDL+ plan for Π and let H be the plan projection of π_t ; π_t is said to be a valid plan for Π iff $H_S(T_m) \neq G$ and the sequence of actions $H_A(T)$ is applicable in $H_S(T)$ for each T in T_H .*

Having clarified the continuous semantics for PDDL+, we proceed by adapting the proposed definitions to the discrete case. In the following, as a convention, we distinguish a history generated according to continuous and discrete semantics using H and H , respectively. Furthermore, we use $\delta \geq Q_{>0}$ (where $Q_{>0} = \{q \geq Q \mid q > 0\}$) as discretisation step.

The following definition of the discrete STP inherits Conditions 1–4 from Definition 9, while Condition 5 is reshaped for detecting when a process has started or stopped in the discrete setting.

Definition 15 (Discrete Significant Time Point). *Let $\delta \geq Q_{>0}$. $T = ht, ni$ is a STP of a history H for Π over $l = [T_s, T_e]$, iff $T \geq l$ and at least one Condition of Definition 9 holds, with Condition 5 reshaped as follows:*

- 5 a process has started (stopped) in T ; formally $H_A(T) = hi$ and $E_{trigg}(T) = hi$ and there exists $\rho \geq P$ such that $H_S(T) \neq pre(\rho)$ ($H_S(T) \neq pre(\rho)$) and there exists a $T^0 = ht^0, n^0i \geq l$ with $t^0 = t - \delta$ such that $H_A(T) = hi$ and $E_{trigg}(T) = hi$ and $H_S(T) \neq pre(\rho)$ ($H_S(T) \neq pre(\rho)$)

Let ξ be the numeric expression that denotes the contribution to the time derivative of some numeric variable, and let $\delta \geq Q_{>0}$, $\Delta(\xi, \delta) = \xi - \delta$ represents the discretisation of ξ according to δ . For example, let $hx, 1.5 - y$ ($\dot{x} = 1.5 - y$) and $\delta = 2$ be a continuous effect and a discretisation parameter, the discretised expression is $\Delta(1.5 - y, \delta) = 3 - y$.

The following definition of the discrete PDDL+ plan projection inherits Rules 1–3 from Definition 13, while Rule 4 is reshaped to enforce how numeric variables change when time advances by a discrete quantity.

Definition 16 (Discrete PDDL+ Plan Projection). *Let $\delta \geq Q_{>0}$, let H be a history for Π over l , let I be an initial state and let π_t be a PDDL+ plan for Π . We say that H is a discrete projection of π_t which starts in I iff H induces the STPs $T_H = hI_0 = ht_s, 0i, \dots, T_m = ht_e, n_mi$ where either $t_{i+1} = t_i + \delta$ or $t_{i+1} = t_i$, and all rules as for Definition 13 apply, except for R4 that becomes:*




Figure 1: History H of the plan $\pi_t = h\pi = h\langle accelerate, 0i, h\langle accelerate, 30i, h\langle t_s = 0, t_e = 90i \rangle \rangle$ for a variant of the CAR domain (Fox & Long, 2006). The history generates 7 STPs, i.e., $T_{\pi} = hT_s = T_0, T_1, T_2, T_3, T_4, T_5, T_e = T_6i$, yielding a piece-wise function consisting of three different monotonous intervals, i.e., $(0, 30)$, $(30, 50)$ and $(50, 90)$. The execution of the action *accelerate* in $T_0 = h\langle 0, 0i$ activates the process *moving* in T_1 ($\mathcal{C}(T_1) = f\langle moving \rangle$) and then the car increases its speed linearly in $(0, 30)$ (according to $\dot{v} = a = 1$ and consequently to $v(t) = t$). In $(30, 0)$, after the second acceleration is performed in $T_2 = h\langle 30, 0i$, the car increases its speed linearly but with a greater slope (according to $\dot{v} = a = 2$ and consequently to $v(t) = 2 \cdot t - 30$). In $T_4 = h\langle 50, 0i$ an event is triggered, i.e., the event which models the explosion of the engine when a critical velocity of 70 is reached. The event resets the acceleration activating the process *drag* in $T_5 = h\langle 50, 1i$ ($\mathcal{C}(T_5) = f\langle moving, drag \rangle$) and deactivating the effect of *moving* on v , which in turns makes the speed decrease exponentially in the interval $(50, 90)$ due to the drag ($\dot{v} = -\frac{1}{5} v$ and then $v(t) = 70 \cdot e^{-\frac{1}{5}(t-50)}$).

R4 for each pair of contiguous STPs $T_i = \langle t_i, n_i \rangle$, $T_{i+1} = \langle t_{i+1}, 0 \rangle$ such that $t_{i+1} = t_i + \delta$, the value of each numeric variable $x \in X$ is updated as:

$$H_S(T_{i+1})[x] = H_S(T_i)[x] + \sum_{\substack{hx^0: i \geq e(\cdot); x^0=x \\ \exists P \text{ such that } H_S(T_i) \neq \text{pre}(\cdot)}} H_S(T_i)[\Delta(\xi, \delta)]$$

and values of unaffected variables remain unchanged (frame-axiom).

Note that plans featuring actions at non-discrete time points do not admit any projection, and are therefore ill-defined under discrete interpretation.⁴ Furthermore, we remark that the discretisation step $\delta \in \mathbb{Q}_{>0}$ is a given parameter (see discussion in Section 7).

Definition 17 (Discrete Valid PDDL+ Plan). *Let π_t be a PDDL+ plan and let H be the plan discrete projection of π_t for $\delta \in \mathbb{Q}_{>0}$; π_t is said to be a valid plan for Π under δ discretisation iff $H_S(T_m) \neq G$ and the sequence of actions $H_A(T)$ is applicable in $H_S(T)$ for each T in T_H .*

2.1 Example: The OVERTAKING-CAR Problem

In order to explain our approach, in the remaining of this paper we will make use of a variant of the well-known LINEAR-CAR PDDL+ problem (Fox & Long, 2006), namely OVERTAKING-CAR.

Example 2.3 (OVERTAKING-CAR Problem). *The OVERTAKING-CAR domain models a number of cars that can move across two lanes, a fast and a slow lane. The position of each car changes over time when the speed is different from 0; the speed in turn changes for the effect of the acceleration. Each car needs to reach a given distance from the origin while avoiding collisions with other cars; to do so a car can switch from one lane to the other. Let cars denote a finite set of vehicles, a PDDL+ planning problem $\Pi = \langle F, X, I, G, A, E, P \rangle$ of OVERTAKING-CAR can be defined as the following:*

- $F = \{ \text{crashed} \} \cup \{ \text{ot}_{car}, \text{eng-on}_{car} \}$, where
 - *crashed* models whether a collision among any vehicle has occurred or not. When the variable is set to true, the goal cannot be achieved anymore;
 - *ot_{car}* models whether a “car” is or not in the fast lane. If *ot_{car}* holds the car is overtaking; otherwise it means that the car is moving on the slow lane;
 - *eng-on_{car}* models whether the engine is on or off;
- $X = \{ \text{d}_{car}, \text{v}_{car}, \text{a}_{car} \}$, where the variables *d_{car}*, *v_{car}* and *a_{car}* express the distance from the origin, the velocity and the acceleration of “car”, respectively;

4. We remark that the plan discrete projection does not model any transient state between two consecutive discrete time points.

- $I = \bigcup_{car \in cars} fh_{car} := D_{car}^I, hv_{car} := 0i, ha_{car} := 0ig$, where $D_{car}^I \in \mathbb{Q}$; for each $car \in cars$, the speed and the distance are specified in the initial state; by closed world assumption the initial state sets to off all engines' cars, constrains all cars to be positioned in the slow lane and crashed is set to false;
- $G = : crashed \wedge \bigwedge_{car \in cars} hv_{car} = 0i \wedge hd_{car} = D_{car}^G i \wedge : eng-on_{car}$, where $D_{car}^G \in \mathbb{Q}$; G requires that (i) all cars have stopped (ii) no collisions have occurred (iii) all cars are at a given distance D_{car}^G from the origin;
- $A = \bigcup_{car \in cars} faccelerate_{car}, fdecelerate_{car}, fovertake_{car}, freturn_{car}, fstart_{car}, fstop_{car}g$ where, for each $car \in cars$, such actions are defined as:
 - $accelerate_{car} = heng-on_{car} \wedge ha_{car} < A_{max}i, fhincrease, a_{car}, 1ig$;
 - $decelerate_{car} = heng-on_{car} \wedge ha_{car} > A_{min}i, fhdecrease, a_{car}, 1ig$;
 - $overtake_{car} = heng-on_{car} \wedge hv_{car} > 0i \wedge : ot_{car}, fot_{car}gi$;
 - $return_{car} = heng-on_{car} \wedge hv_{car} > 0i \wedge ot_{car}, f: ot_{car}gi$;
 - $start_{car} = h: eng-on_{car}, feng-on_{car}gi$;
 - $stop_{car} = heng-on \wedge hv_{car} = 0i \wedge ha_{car} = 0i, f: eng-on_{car}gi$.

The actions $overtake_{car}$ and $return_{car}$, are used to model the movement of “car” from the slow lane to the fast lane and vice versa, respectively.

- $E = \bigcup_{\substack{car, car^0 \in cars \\ car \neq car^0}} fcrash-fast_{car, car^0}, fcrash-slow_{car, car^0}g$; these events model the possible collisions that can occur between two moving cars, i.e., car and car^0 , if those are moving in the same lane at a distance less than a critical distance $D_T \in \mathbb{Q}_{>0}$. They are defined as follows:
 - $crash-fast_{car, car^0} = hhd_{car} \quad d_{car^0} < D_T i \wedge hd_{car} \quad d_{car^0} \quad 0i \wedge ot_{car} \wedge ot_{car^0} \wedge : crashed, fcrashedgi$;
 - $crash-slow_{car, car^0} = hhd_{car} \quad d_{car^0} < D_T i \wedge hd_{car} \quad d_{car^0} \quad 0i \wedge : ot_{car} \wedge : ot_{car^0} \wedge : crashed, fcrashedgi$;
- $P = \bigcup_{car \in cars} fmoving_{car}g$ where, for each $car \in cars$, such processes are defined as:
 - $moving_{car} = heng-on_{car}, fhv_{car}, a_{car}i, hd_{car}, v_{car}ig$, this process models the dynamics of the car: the time derivative of the displacement is the speed (i.e., $\dot{d}_{car} = v_{car}$), whereas the time derivative of the speed is the acceleration (i.e., $\dot{v}_{car} = a_{car}$);

The numeric constants involved in the definition of I and G are specified such that:

- for each pair of cars, i.e., $car, car^0 \in cars$ with $car \neq car^0$, then $jD_{car}^I \quad D_{car^0}^I j < D_T$; this constraint prevents a collision event from being triggered in the initial state

- for each car 2 cars then $D_{car}^G > D_{car}^I$;
- for each pair of cars, i.e., car, car^0 2 cars with $car \notin car^0$ and then $D_{car}^G < D_{car^0}^G$ ($D_{car}^G > D_{car^0}^G$); this implies that to achieve the goal, it is necessary to use the fast lane;
- the maximum acceleration and deceleration are bounded by two quantities, i.e., A_{max} and A_{min} , respectively.

In OVERTAKING-CAR, the actions to be sought and scheduled are the accelerations and decelerations to reach the goal, while the actions to handle overtaking are used to avoid collisions. Accelerate and decelerate actions have an impact on the speed and therefore on the actual positions of the cars. Such aspects are directly modelled through processes. Each process indeed models how the distance from the origin is affected; this is given as a function of speed that in turn is controlled by the current acceleration.

3. Translating PDDL+ to PDDL2.1

PDDL+ problems differ from PDDL2.1 ones for the presence of processes and events. In this and the following sections, we devise translations that transform all processes and events into regular actions enriched with additional predicates so that every valid plan that is found in the process and event-free translation, retains its validity on the discretised version of the PDDL+ problem it has been generated from. We do so by explicitly formulating a simulation action that lets the planning engine wait and observe the state of the world for a given amount of time. Instead, when the planner picks some action, time does not flow; rather the state is instantaneously modified by the planning engine.

To make this operational, there are a number of challenges to pursue: (i) we need to capture what processes are active in a given state so that the time-discretised continuous update of the state consistently reflects what the dynamical specification of the system prescribes; (ii) we need to take care of the potentially complex cascade of events that may be triggered for each encountered state.

We proceed in a modular fashion. In what follows we firstly present a solution to point (i) with a (straightforward) exponential encoding, and then with a more sophisticated polynomial translation. Both are guaranteed to work for event-free PDDL+ problems. Then we face point (ii) by showing how also events can be translated into actions with conditional effects with a translation step that is modular to how we tackle point (i).

3.1 Exponential Translation

Given an event-free PDDL+ problem $\Pi = \langle hF, X, I, G, A, \cdot, P \rangle$, we define a discrete context C , hereinafter simply referred to as context, to be a non-empty subset of processes, and denote with $P^+(P)$ the set of non-empty subsets of P , that is the set of all possible contexts.

For an event-free PDDL+ problem Π , the exponential translation generates a PDDL2.1 problem $\Pi_{EXP} = \langle hF, X, I, G, A \cup \{SIM\}, \delta \rangle$, discretised in δ . Π_{EXP} is almost identical to Π but for the absence of processes and the presence of the special action *SIM* playing the role of the simulator, i.e., what changes when time goes forward. *SIM* is defined as follows:

$$\begin{aligned}
 pre(SIM) &= > \\
 eff(SIM) &= \bigcup_{C \in 2P^+(P)} fcontpre(C) \triangleright conteff(C)g
 \end{aligned}$$

where

$$\begin{aligned}
 contpre(C) &= \bigwedge_{2PnC} : pre(\rho) \wedge \bigwedge_{2P \setminus C} pre(\rho) \\
 conteff(C) &= \bigcup_{x \in X} fhinc, x, \sum_{\substack{hx^\delta; i \in 2e(\cdot); \\ x^\delta = x}} \Delta(\xi, \delta) ig
 \end{aligned}$$

Intuitively, the action SIM organises all possible contexts within a unique action, delegating to each conditional effect (i) the conditions under which a context is triggered and (ii) the consequences that such a context has on the state after some time δ has passed. Point (i) is formalised by conjoining two conjunctions: the first ensures that no other process of some other context has its precondition satisfied ($\bigwedge_{2PnC} : pre(\rho)$); the second ensures that all the preconditions of a given context are satisfied ($\bigwedge_{2P \setminus C} pre(\rho)$). Let x be some numeric variable of our problem, point (ii) is obtained by summing the contribution of each process within the context.

Ultimately, we reflect the effect of the time-passing action on the overall makespan of the plan directly in the cost function of our problem, differentiating therefore whether the planning engine chooses some action, or lets time go for some δ time: $c(a) = 0$ if $a \in A$ while $c(a) = \delta$ if $a = SIM$.

Example 3.1 (EXP Translation - Continuing on Example 2.3). *Let $\Pi = \langle hF, X, I, G, A, E, P \rangle$ be an OVERTAKING-CAR planning instance for the domain reported in Example 2.3 in which $cars = \langle fcar1, car2 \rangle g$. In this example, we compile Π omitting the events (therefore as if it were $E = \langle \cdot \rangle$); events will be dealt with later in the paper. The components of the tuple Π are therefore:*

$$\begin{aligned}
 F &= \langle feng-on_{car1}, eng-on_{car2}, ot_{car1}, ot_{car2}, crashed \rangle g \\
 X &= \langle fd_{car1}, d_{car2}, v_{car1}, v_{car2}, a_{car1}, a_{car2} \rangle g; \\
 I &= \langle fh_{car1} := D_{car1}^I i, hd_{car2} := D_{car2}^I i, hv_{car1} := 0i, hv_{car2} := 0i, ha_{car1} := 0i, ha_{car2} := 0i \rangle g; \\
 G &= \langle crashed \wedge hd_{car1} = D_{car1}^G i \wedge hd_{car2} = D_{car2}^G i \wedge hv_{car1} = 0i \wedge hv_{car2} = 0i \wedge : eng-on_{car1} \wedge \\
 &\quad : eng-on_{car2}; \\
 A &= \langle faccelerate_{car1}, accelerate_{car2}, overtake_{car1}, overtake_{car2}, return_{car1}, return_{car2}, \\
 &\quad start_{car1}, start_{car2}, stop_{car1}, stop_{car2} \rangle g; \\
 E &= \langle \cdot \rangle; \\
 P &= \langle fmoving_{car1}, moving_{car2} \rangle g.
 \end{aligned}$$

The EXP reformulation of Π discretised in δ is $\Pi_{EXP} = \langle hF, X, I, G, A [fSIMg, ci \rangle$ where SIM is the simulation action as for EXP specification. Let W_C be the single conditional effect

of SIM such that $C \supseteq P^+(P)$ is a context, and let “m1” and “m2” be short for processes “moving_{car1}” and “moving_{car2}”. The SIM action is as follows:

$$\begin{aligned} pre(SIM) &= > \\ eff(SIM) &= fW_{fm1g}, W_{fm2g}, W_{fm1;m2g} \end{aligned}$$

where

$$\begin{aligned} W_{fm1g} &= eng-on_{car_1} \wedge : eng-on_{car_2} \triangleright fhinc, d_{car_1}, v_{car_1} \ \delta i, hinc, v_{car_1}, a_{car_1} \ \delta ig \\ W_{fm2g} &= : eng-on_{car_1} \wedge eng-on_{car_2} \triangleright fhinc, d_{car_2}, v_{car_2} \ \delta i, hinc, v_{car_2}, a_{car_2} \ \delta ig \\ W_{fm1;m2g} &= eng-on_{car_1} \wedge eng-on_{car_2} \triangleright fhinc, d_{car_1}, v_{car_1} \ \delta i, hinc, v_{car_1}, a_{car_1} \ \delta i, \\ &\quad hinc, d_{car_2}, v_{car_2} \ \delta i, hinc, v_{car_2}, a_{car_2} \ \delta ig. \end{aligned}$$

In the following lemma, we show the soundness and completeness properties of the EXP translation for a PDDL+ event-free problem. By completeness and soundness, we mean that if Π admits a valid solution then the corresponding problem Π_{EXP} admits one solution too, and vice-versa.

Lemma 1 (Soundness and Completeness of EXP for an Event-free PDDL+ Problem). *Let $\Pi = hF, X, I, G, A, ;, Pi$ be a PDDL+ problem, and let $\Pi_{EXP} = hF, X, I, G, A [fSIMg, ci$ be the PDDL2.1 problem obtained by using the EXP translation discretised in δ . Π admits a solution under δ discretisation iff so does Π_{EXP} .*

Proof. () Let $\pi_t = h\pi, h0, t_e i i$ be a valid solution for Π (assume w.l.o.g. $t_s = 0$) under δ discretisation and let π_{EXP} be a PDDL2.1 plan constructed in such a way that: (i) for each ha, ti in π then a^δ is in π_{EXP} (where a^δ is the compiled version of a); (ii) for each $ha_j, t_j i, ha_j, t_j i$ with $a_j \ a_j$ in π then $a_j^\delta \ a_j^\delta$ holds in π_{EXP} and (iii) a sequence, possibly empty, of SIM actions has to be placed before each action $a_j^\delta \supseteq \pi_{EXP}$ and at the end of π_{EXP} according to the following structure

$$\pi_{EXP} = hhSIM i \ \frac{t_0}{\delta}, a_0^\delta, hSIM i \ \frac{t_1 - t_0}{\delta}, \dots, a_{n-1}^\delta, hSIM i \ \frac{t_e - t_{n-1}}{\delta} i$$

where $hSIM i \ k$, with $k \supseteq \mathbb{N}$, indicates k repetitions of SIM .

In order to prove that π_{EXP} is a valid solution for Π_{EXP} , it suffices to show that, let $\tau = hH_S(T_0), \dots, H_S(T_m) i$ be the sequence of states associated to each STP of H , and $\tau^\delta = hs_0, \dots, s_m i$ be the sequence of states generated by iteratively executing π_{EXP} , $H_S(T_i)$ and s_i are equivalent (agree on all values for $F [X$) for each $i \supseteq [0..m]$. We prove this by induction on τ (τ^δ).

Firstly, we have to show how $n = j\pi j$ and $m = j\tau j - 1$ are related and then that $j\tau j = j\tau^\delta j$ holds. Given π , according to R1 of Definition 16, we have n instantaneous transitions and, according to R4, we have $\frac{t_e - t_s}{\delta} = \frac{t_e}{\delta}$ temporal transitions, i.e., the number of discrete advancements of time. It follows that $m = n + \frac{t_e - t_s}{\delta}$. As shown in π_{EXP} definition, there is an action $a_i^\delta \supseteq \pi_{EXP}$ for each action $a_i \supseteq \pi$ with $i \supseteq [0..n - 1]$, and a SIM action for each

discrete advancement of time, i.e., $\frac{t_e}{j}$. Then $j\pi_{\text{EXP}}j = n + \frac{t_e}{j} = m$. Then π_{EXP} induces m transitions from I , i.e., $j\tau^0j = m + 1$, and finally we got that $j\tau j = j\tau^0j$.

The base case ($i = 0$) trivially proves true as $H_S(T_0) = I$ and $s_0 = I$. For the induction step, we assume truly the statement for some $i < j\tau j$, and prove this for $i + 1$ by considering the two types of transitions occurring between two contiguous STPs in H .

Instantaneous transition. Let $T_i = ht_i, n_i i$ and $T_{i+1} = ht_i, n_i + 1 i$ be two STPs of H . Rule 2 of Definitions 13-16 implies that $H_A = ha_i i \notin hi$. Since π_t is a valid solution for Π , we know that $H_S(T_i) \neq pre(a_i)$ therefore, by inductive hypothesis, $s_i \neq pre(a_i^0)$. Since a_i and a_i^0 are the same operator, it is easy to see that the outcomes of the transitions $\gamma(H_S(T_i), a_i)$ and $\gamma(s, a_i^0)$ are equivalent.

Temporal transition. Let i be an index such that $T_i = ht_i, n_i i$ and $T_{i+1} = ht_i + \delta, 0 i$ are two STPs of H .

Note that the *SIM* action features a conditional effect for each possible context $C \subseteq P^+(P)$. By definition of $P^+(P)$, contexts differ from one another for at least one process. It follows that, by definition of *contpre*(\cdot), for each $C, C^0 \subseteq P^+$ with $C \neq C^0$, *contpre*(C) and *contpre*(C^0) are mutually exclusive. Therefore, for each application of the *SIM* action, at most one conditional effect is activated.

We denote with *SIM* $_i$ the application of *SIM* in the i -th state of τ^0 . State s_i induces the context $C(s_i) = f\rho \subseteq P$, $s_i \neq pre(\rho)g$ and then *SIM* $_i = h>, contpre(C(s_i)) \triangleright conteff(C(s_i))i = hcontpre(C(s_i), conteff(C(s_i))i = h>, conteff(C(s_i))i$; indeed, we can remove each conditional effect that does not hold in s_i . To show that the state produced by the application of *SIM* $_i$ in s_i is equivalent to that produced according to the semantics of the discrete projection of π_t , when a quantum of δ time passes, it is sufficient to focus on a single variable $x \subseteq X$ that is affected by the process, and then generalise to all the numeric variables.⁵ In particular, action *SIM* $_i$ modifies x according to:

$$hinc, x, \sum_{\substack{hx^0; i2e(\cdot); x^0=x \\ 2C(s_i)}} \Delta(\xi, \delta) i$$

which corresponds to enforcing the transition as follows:

$$s_{i+1}[x] = s_i[x] + \sum_{\substack{hx^0; i2e(\cdot); x^0=x \\ 2C(s_i)}} s_i[\Delta(\xi, \delta)] \quad (1)$$

According to Rule 4 of Definition 16, each numeric variable $x \subseteq X$ changes over δ by using the active processes in $C(H_S(T_i)) = f\rho \subseteq P$, $H_S(T_i) \neq pre(\rho)g$ according to:

$$H_S(T_{i+1})[x] = H_S(T_i)[x] + \sum_{\substack{hx^0; i2e(\cdot); x^0=x \\ 2C(H_S(T_i))}} H_S(T_i)[\Delta(\xi, \delta)] \quad (2)$$

Since $H_S(T_i)$ is equivalent to s_i , by the inductive hypothesis, then the induced contexts are the same, i.e., $C(s_i) = C(H_S(T_i))$. It follows that the right-hand side expressions of Formulae 1-2 are equivalent. Thus, $H_S(T_{i+1})$ and $s_{i+1} = \gamma(s_i, SIM)$ are equivalent.

5. Note that, the study of the equivalence of s_{i+1} and $H_S(T_{i+1})$ can be circumscribed to studying the numeric variables X and ignoring the propositional ones, since the P processes of can only affect numeric variables. Same for the conditional effects of *SIM* they are built from.

any process starts operating. This lets the planning engine safely evaluate all variables each process effect depends on. In order to achieve this last point, each numeric effect in some process is rewritten before being transformed into an action. The rewriting manipulates each formula in a way to substitute every occurrence of a variable from X with its compilation image in X^{cp} .⁶ Let ξ be a formula, we denote with $\sigma(\xi, X^{cp})$ the result of such a rewriting.

This translation trades the exponential blow-up caused by the context-switching operation of EXP, with an increase in the length of the plan. The resulting formulation is sound, complete and, more interestingly, is polynomial on the size of the input. For this reason, we call this translation POLY and detail in what follows its precise definition and functioning.

Let $\Pi = \langle hF, X, I, G, A, \cdot, Pi \rangle$ and $\delta \geq 0$ be an event-free PDDL+ problem and a discretisation parameter. POLY generates a new PDDL2.1 problem $\Pi_{POLY} = \langle hF [D [fpauseg, X [X^{cp}, I, G \wedge : pause, A_c [A_P [fstart, endg, ci \rangle$ such that:

$$\begin{aligned}
 X^{cp} &= fx^{copy} \ j \ x \ 2 \ Xg \\
 D &= \bigcup_{\substack{ne \in e(\cdot) \\ 2P}} fdone_{ne}g \\
 A_c &= fhpre(a) \wedge : pause, eff(a) \ i \ j \ a \ 2 \ Ag \\
 start &= h: pause, fpauseg [\bigcup_{x \in X} fhasgn, x^{copy}, x \ i \ g \ i \\
 end &= h \bigwedge_{done \in D} done \wedge pause, f: pauseg [\bigcup_{done \in D} f: doneg \ i \\
 A_P &= \bigcup_{\substack{ne: hx; \ i \ 2 \ e(\cdot) \\ 2P}} fhpause \wedge : done_{ne}, f\sigma(pre(\rho), X^{cp}) \triangleright fhinc, x, \Delta(\delta, \sigma(\xi, X^{cp})) \ i \ g \ g [fdone_{ne}g \ i \ g
 \end{aligned}$$

As it is possible to observe, at any discrete time step, the planning engine can decide to let time pass by an amount of δ and does so by deciding to execute action *start*. From that moment onward, no action from A_c can be executed, and only when all conditional process effects are applied, the planning engine can come back into actual planning mode ($: pause$). A_P encompasses all such processes effects, and delegates to a conditional effect the check and the consequent update of the variables according to their past value ($\sigma(pre(\rho), X^{cp})$ for the precondition of the process, and $\sigma(\xi, X^{cp})$ for the right-hand side of the numeric effect) under the proper Δ discretisation. Note that, action *start* also ensures that all the variables are copied through an assignment operation, which is responsible for iterating over all numeric variables of the problem and updating their value for the next round of simulation (the snippet $\bigcup_{x \in X} fhasgn, x^{copy}, x \ i \ g$).

Similarly to the exponential translation, also in this case we make the planning engine aware of the passage of time through the cost function. As, however, we do not have only one action that reflects such a passage of time, we attribute a non-zero cost only to action *start*. That is: $c(a) = \delta$ if $a = start$, 0 otherwise.

Example 3.2 (POLY Translation - Continuing on Example 2.3). *Let $\Pi = \langle hF, X, I, G, A, E, Pi \rangle$ be an OVERTAKING-CAR planning instance for the domain reported in Example 2.3 in which $cars = fcar1, car2g$. Again, let us assume for now that $E = \cdot$.*

6. Recall that every numeric expression is a formula.

The process labels are shortened using the same convention provided in Example 3.1. However, if a process admits more than one continuous numeric effect, another convention is introduced. For example, consider the grounded process moving_{car1} having two numeric continuous effects, i.e., $\text{hd}_{car1, v_{car1}i}$ and $\text{hv}_{car1, a_{car1}i}$. Such effects are labelled as “ m_{d1} ” and “ m_{v1} ”, respectively. That is, “ m_{d1} ” is the numeric continuous effects of moving_{car1} affecting d_1 whereas “ m_{v1} ” is the numeric continuous effects of moving_{car1} affecting v_1 .

The polynomial reformulation of Π discretised in δ is $\Pi_{\text{POLY}} = \text{hF} [D [\text{fpauseg}, X [X^{\text{cp}}, I, G [f: \text{pauseg}, A_c [A_P [\text{fstart}, \text{endgi}$ where:

$$\begin{aligned}
 D &= \bigcup_{\substack{\text{ne2e} () \\ \text{2fm1,m2g}}} \text{fdone}_{\text{ne}g} = \text{fdone}_{m_{d1}}, \text{done}_{m_{v1}}, \text{done}_{m_{d2}}, \text{done}_{m_{v2}}, g \\
 X^{\text{cp}} &= \bigcup_{\text{car2fcar1:car2g}} \text{fd}_{\text{car}}^{\text{copy}}, \text{v}_{\text{car}}^{\text{copy}}, \text{a}_{\text{car}}^{\text{copy}} g = \text{fd}_{\text{car1}}^{\text{copy}}, \text{d}_{\text{car2}}^{\text{copy}}, \text{v}_{\text{car1}}^{\text{copy}}, \text{v}_{\text{car2}}^{\text{copy}}, \text{a}_{\text{car1}}^{\text{copy}}, \text{a}_{\text{car2}}^{\text{copy}} g \\
 \text{start} &= \text{h: pause, fpauseg} [\text{fhasgn}, \text{d}_{\text{car1}}^{\text{copy}}, \text{d}_{\text{car1}i}, \text{hasgn}, \text{d}_{\text{car2}}^{\text{copy}}, \text{d}_{\text{car2}i}, \text{hasgn}, \text{v}_{\text{car1}}^{\text{copy}}, \text{v}_{\text{car1}i}, \\
 &\quad \text{hasgn}, \text{v}_{\text{car2}}^{\text{copy}}, \text{v}_{\text{car2}i}, \text{hasgn}, \text{a}_{\text{car1}}^{\text{copy}}, \text{a}_{\text{car1}i}, \text{hasgn}, \text{a}_{\text{car2}}^{\text{copy}}, \text{a}_{\text{car2}i} g \\
 \text{end} &= \text{hdone}_{m_{d1}} \wedge \text{done}_{m_{v1}} \wedge \text{done}_{m_{d2}} \wedge \text{done}_{m_{v2}} \wedge \text{pause}, \\
 &\quad f: \text{done}_{m_{d1}} \wedge : \text{done}_{m_{v1}} \wedge : \text{done}_{m_{d2}} \wedge : \text{done}_{m_{v2}} \wedge : \text{pausegi} \\
 A_P &= \text{fa}_{m_{d1}}, \text{a}_{m_{v1}}, \text{a}_{m_{d2}}, \text{a}_{m_{v2}} g \\
 a_{m_{d1}} &= \text{hpause} \wedge : \text{done}_{m_{d1}}, \text{fengine-on}_{\text{car1}} \triangleright \text{fhinc}, \text{d}_{\text{car1}}, \text{v}_{\text{car1}}^{\text{copy}} \delta \text{igg} [\text{fdone}_{m_{v1}} g \\
 a_{m_{v1}} &= \text{hpause} \wedge : \text{done}_{m_{v1}}, \text{fengine-on}_{\text{car1}} \triangleright \text{fhinc}, \text{v}_{\text{car1}}, \text{a}_{\text{car1}}^{\text{copy}} \delta \text{igg} [\text{fdone}_{m_{d1}} g \\
 a_{m_{d2}} &= \text{hpause} \wedge : \text{done}_{m_{d2}}, \text{fengine-on}_{\text{car2}} \triangleright \text{fhinc}, \text{d}_{\text{car2}}, \text{v}_{\text{car2}}^{\text{copy}} \delta \text{igg} [\text{fdone}_{m_{d2}} g \\
 a_{m_{v2}} &= \text{hpause} \wedge : \text{done}_{m_{v2}}, \text{fengine-on}_{\text{car2}} \triangleright \text{fhinc}, \text{v}_{\text{car2}}, \text{a}_{\text{car2}}^{\text{copy}} \delta \text{igg} [\text{fdone}_{m_{v2}} g
 \end{aligned}$$

Note that the use of X^{cp} variables is crucial to make deterministic the outcome of the time-passage simulation sequence, i.e., $\text{hstart}, \text{seq}(A_P), \text{end}i$, regardless of the chosen ordering in $\text{seq}(A_P)$.

In the following lemma, we show that POLY is sound and complete for a PDDL+ event-free problem.

Lemma 2 (Soundness and Completeness of POLY for an Event-free PDDL+ Problem). *Let $\Pi = \text{hF}, X, I, G, A, \cdot, Pi$ be a PDDL+ problem, and let $\Pi_{\text{POLY}} = \text{hF} [D [\text{fpauseg}, X [X^{\text{cp}}, I, G [f: \text{pauseg}, A_c [A_P [\text{fstart}, \text{endgi}$ be the PDDL2.1 problem obtained by using the POLY translation discretised in δ . Π admits a solution under δ discretisation iff so does Π_{POLY} .*

Proof. () Let $\pi_t = \text{h}\pi, \text{h}0, t_ei$ be a valid solution for Π (assume w.l.o.g. $t_s = 0$) under δ discretisation, and let π_{POLY} be a PDDL2.1 plan constructed in such a way that: (i) for each ha, ti in π then a^0 is in π_{POLY} (where a^0 is the compiled version of a); (ii) for each $\text{ha}_i, t_i, \text{ha}_j, t_ji$ with $a_i \text{ } a_j$ in π then $a_i^0 \text{ } a_j^0$ holds in π_{POLY} (iii) a sequence, possibly empty, of sequences having the form $\text{wait} = \text{hstart}, \text{seq}(A_P), \text{end}i$ (where $\text{seq}(A_P)$ is any sequencing of all A_P operators) has to be placed before each action a_i^0 in π_{POLY} and at the end of π_{POLY} according to the following structure:

$$\pi_{\text{POLY}} = hhwaiti \frac{t_0}{\delta}, a_0^0, hwaiti \frac{t_1}{\delta}, \dots, a_{n-1}^0, hwaiti \frac{t_n}{\delta} i$$

Much as we do for Lemma 1, we proceed by induction over the states τ and τ^θ , noticing that τ^θ can be constructed using *wait* as if it was a single transition, therefore leading us to have $j\tau j = j\tau^\theta j$. Differently from Lemma 1, Π and Π_{POLY} have different variables, i.e., the set of variables of Π_{POLY} contains the set of variables of Π . Hence, given two elements of τ and τ^θ , i.e., s_i and $H_S(T_i)$, we say that they are equivalent if they match over the variables of Π , i.e., $F \llbracket X$ and $H_S(T_i) \not\equiv : \text{pause}$.

The base case ($i = 0$) trivially proves true as $H_S(T_0) = I \not\equiv : \text{pause}$ and $s_0 = I$. The *instantaneous transition* is similar to what has been discussed in Lemma 1 since a_i and its counterpart a_i^θ are the same operator except for the *: pause* precondition of a_i^θ .

The non-trivial aspect to prove is the inductive case when there is a *temporal transition*. Starting from two equivalent states in τ and τ^θ , i.e., $H_S(T_i)$ and s_i , we need to prove that $H_S(T_{i+1})$, with $T_i = ht_i, n_i i$ and $T_{i+1} = ht_i + \delta, 0 i$, is equivalent to $s_{i+1} = \gamma(s_i, \text{wait})$. First, we discuss the applicability of the sequence *wait*. The first action *start* has *: pause* as a precondition, so, since the inductive hypothesis holds, it is applicable in s_i making *pause* true; after *start* has been applied, the only executable operators are those belonging to A_P . Then, any sequencing of A_P generates a state in which $\bigwedge_{done2D} done$ holds, thus making *end* applicable, too.

Now we proceed to show that the two compared states are equivalent under variables $F \llbracket X$. Note that each operator in A_P only modifies the X variables. Thus both variables in X^{cp} and in F persist during and after the application of any sequencing of A_P . The only thing left to prove is to show the equivalence of the states w.r.t. variables from X .

Recall that the *start* operator makes a copy of all the X variables in the corresponding duplicates in X^{cp} , therefore the set of active processes does not change starting from s_i . Indeed, all compiled processes' preconditions only involve variables from X^{cp} .

Each operator $a = hpause \wedge : done_{ne}, fc \triangleright e, done_{gi} \geq A_P$ has a conditional effect that refers to a process ρ , i.e., $c \triangleright e$ with $c = \sigma(\text{pre}(\rho), X^{cp})$ and $e = hinc, x, \Delta(\delta, \sigma(\xi, X^{cp})) i$, which increases x if a formula defined over X^{cp} , i.e., $\sigma(\text{pre}(\rho), X^{cp})$, holds in the state where a is applied. Note moreover that the activation of the conditional effects is independent of the order in which actions from A_P are evaluated. In fact, the preconditions of the conditional effects of A_P are expressed in X^{cp} variables which, as noticed above, remain constant throughout the execution of A_P .

To see why $H_S(T_{i+1})$ and s_{i+1} are equivalent, note that, given a variable $x \geq X$, the effects on x is given by a subset of A_P , i.e., $A_P(x)$, defined as the following:

$$A_P(x) = \bigcup_{\substack{ne:hx^\theta; i2e () \\ x^\theta=x; 2P;}} fhpause \wedge : done_{ne}, f\sigma(\text{pre}(\rho), X^{cp}) \triangleright fhinc, x^\theta, \Delta(\sigma(\xi, X^{cp}), \delta)) igg \llbracket fdone_{ne} g i g$$

which can be redefined in $A_P(x, s_i)$ by substituting the right-hand side of the numeric effects with all the variables from s_i :

$$A_P(x, s_i) = \bigcup_{\substack{ne:hx^\theta; i2e () \\ x^\theta=x; 2P}} fhpause \wedge : done_{ne}, f\sigma(\text{pre}(\rho), X^{cp}) \triangleright fhinc, x^\theta, s_i[\Delta(\xi, \delta)] igg \llbracket fdone_{ne} g i g$$

Now, observe that the conditional effects modifying the value of x during *wait* are those created from process ρ where $s_i \neq \text{pre}(\rho)$, i.e., those processes belonging to the context $\mathcal{C}(s_i)$. So we can express $s_{i+1}[x]$ as the following summation:

$$s_{i+1}[x] = \gamma(s_i, \text{wait})[x] = s_i[x] + \sum_{\substack{hx^\delta; i2e(\cdot); x^\delta=x \\ 2\mathcal{C}(s_i)}} s_i[\Delta(\xi, \delta)] \quad (3)$$

According to Rule 4 of Definition 16, each numeric variable $x \in X$ changes over δ by the active processes in $\mathcal{C}(H_S(T_i)) = f\rho \in P, H_S(T_i) \neq \text{pre}(\rho)g$:

$$H_S(T_{i+1})[x] = H_S(T_i)[x] + \sum_{\substack{hx^\delta; i2e(\cdot); x^\delta=x \\ 2\mathcal{C}(H_S(T_i))}} H_S(T_i)[\Delta(\xi, \delta)] \quad (4)$$

Since $H_S(T_i)$ is equivalent to s_i , by the inductive hypothesis, then the induced contexts are the same, i.e., $\mathcal{C}(s_i) = \mathcal{C}(H_S(T_i))$. It follows that the right-hand side expressions of Formulae 3-4 are equivalent. Thus, $H_S(T_{i+1})$ and $s_{i+1} = \gamma(s_i, \text{wait})$ are equivalent.

Therefore, by induction, $s_i = H_S(T_i)$ for all i .

(\cdot) Note that every plan solving Π_{POLY} is structured alternating an agent's action and a sequence (possibly empty) of *wait* sequences. Let $\pi_{\text{POLY}} = h\text{wait } t_0, a_0^l, \text{wait } t_1 \dots t_0, \dots, a_{n-1}^l, \text{wait } t_n \dots t_{n-1} i$ be such a plan. The mapping from π_{POLY} to $\pi_t = h\pi, h0, t_e i$ is similar to what was done for Lemma 1. All the occurrences of the sequence *wait* in π_{POLY} have to be ignored in building π . Starting from π_{POLY} we can build a valid PDDL+ plan $\pi_t = h\pi, h0, t_e i$ as follows: (i) for each action a_i^l in π_{POLY} such that $a_i^l \notin \text{start, endg} [A_P]$, then $ha_i, t_i i$ is in π , where t_i is equal to δ multiplied for the occurrences of *start* in π_{POLY} before a_i^l ; (ii) for each a_i^l, a_j^l such that $a_i^l \prec a_j^l$ in π_{POLY} then $ha_i, t_i i \prec ha_j, t_j i$ holds in π and (iii) t_e is equal to δ multiplied by the number of *start* in π_{POLY} . The plan we get by using this transformation is:

$$\pi_t = hh\pi = ha_0, t_0 i, ha_1, t_1 i, \dots, ha_{n-1}, t_{n-1} i, h0, t_e i$$

In order to show the validity of π_t , we reason on the discrete projection H of π_t that is determined as shown in Lemma 1 by using Definition 13. That gives us that the number of STPs of H equates the number of *wait* sequence plus the number of non simulating actions of π_{POLY} , i.e., $|T_H| = j\pi + t_e$, where $j\pi$ is the number of non simulating action in π_{POLY} and t_e is the number of *wait* sequence in π_{POLY} and then $|T_H| = j\pi_{\text{POLY}}|$. It is easy to see that, given that π_{POLY} is valid, and proceeding by reasoning by induction over the induced sequences of states τ and τ^δ as done for the opposite direction, π_t can be proved valid for the problem Π under δ discretisation. \square

3.3 Handling Events

An event in PDDL+ models can be triggered at any time during the execution of a plan, and it is necessary to track whether and how such an event changes the state. More importantly, the semantics of PDDL+ prescribes that a cascade of events may also occur.

In order to handle this behaviour, we devise a new action, namely *SIMEV*, that is responsible for keeping track of the arisen events, and their impacts on the state. This action does so by encoding in one single unit the potentially repeated check and execution of several events via a sophisticated usage of conditional effects that are evaluated in rounds. *SIMEV* makes use of 4 sets of conditional effects:

1. W_{trig} that is responsible for actually updating the state with all events having their precondition satisfied;
2. W_{red} that is responsible for keeping track of whether there is some event triggered;
3. W_{satu} that is responsible for capturing whether no event is triggered in the previous round;
4. $W_{?}$ that captures the situation where there is some inconsistency caused by either a set of mutex events active at the same time or a cyclic sequence of events being triggered.

The formalisation of such conditional effects in PDDL2.1 makes use of a number of additional fresh predicates that are accumulated in set F_E . We have a fact *sim-ev* that signals the beginning of the event simulation; then we have a fact *fired_e* for each event in $\varepsilon \in E$.

The *SIMEV* action is formalised in such a way that:

$$\begin{aligned} pre(SIMEV) &= sim-ev \\ eff(SIMEV) &= W_{trig} [W_{red} [W_{satu} [W_{?} \end{aligned}$$

where:

$$\begin{aligned} W_{trig} &= \bigcup_{\substack{c.e \in E \\ "2E}} fpre(\varepsilon) \wedge c \triangleright eg \\ W_{red} &= \bigcup_{"2E} fpre(\varepsilon) \triangleright fired_{e}.g \\ W_{satu} &= \left\{ \bigwedge_{"2E} (: pre(\varepsilon) _ fired_{e}) \triangleright f: sim-ev.g [\bigcup_{"2E} f: fired_{e}.g \right\} \\ W_{?} &= \left\{ \left(\bigvee_{\substack{": "0 2E: \\ "0 \wedge \\ mutex(";"0)}} pre(\varepsilon) \wedge pre(\varepsilon^0) \right) _ \bigvee_{"2E} pre(\varepsilon) \wedge fired_{e} \triangleright fundefined.g \right\} \end{aligned}$$

As it is possible to observe from the snippet above, *SIMEV* captures which events have been triggered and, on the one hand, applies their effects, and on the other hand, memorises whether at least one event has been executed. If that is the case, the action needs to re-evaluate the conditional effects; indeed, a cascade of events can be triggered. It is easy to see that the triggering of events is blocked whenever the action detects a cycle, i.e., an event that is deemed to be executed more than once. For this reason, *SIMEV* can

be performed up to jEj times, that is, after the termination condition induced by W_{satu} is reached. Observe that effect ($f: sim-evg [\bigcup_{\cdot 2E} f: fired.g$) not only interrupts the execution of the simulation of the events but also resets all *fired* facts; this way, we keep the memory ready for the next round of simulation of events.

The very last set of conditional effects ensures that the reached state does not contain cycles or mutexes events that can be executed at the same time. If either of these two situations arises, *SIMEV* generates an inconsistent state, thereby denoted by the special state *undefined*. This check guarantees to prune states where interfering events leading to non-deterministic outcomes or an infinite cascade of events arise. This complies with the semantics restrictions imposed by previous works in PDDL+ (Shin & Davis, 2005; Fox & Long, 2006; Fox et al., 2005).

Note that if the PDDL+ problem to be discretised satisfies the required assumptions about the events, that is: (i) events have to deactivate themselves, (ii) events have to be fired at most once for a given time and (iii) the PDDL+ task is event-deterministic; then W_{\cdot} is not necessary. Moreover, in some cases, W_{\cdot} may be redundant; suppose to have a state where only a single event ε is triggered, and this does not self-deactivate. In such a case *SIMEV* would have two conflicting conditional effects, i.e., W_{trig} and W_{satu} , which make *fired* true and false at the same time. According to the state transition function, this leads to an undefined state and so W_{\cdot} is redundant. We have however included W_{\cdot} in the conditional effects of *SIMEV*, making the failure to generate a successor state more explicit. This choice makes the encoding more robust w.r.t. violations of the PDDL+ assumptions.

The presented exponential and polynomial translations can be extended to address PDDL+ planning problems with events. This is done by enforcing the action *SIMEV* to be applied in the initial state and switching back to event simulating modality (*sim-ev*) after any occurrence of an instantaneous action (in both translations), and after the execution of the action *SIM* and *end* in EXP and POLY, respectively. Let a be an action, we denote with *ev-check*(a) the new action a^{θ} such that: $pre(a^{\theta}) = pre(a) \wedge sim-ev$ and $eff(a^{\theta}) = eff(a) [fsim-evg$.

We are now ready to summarise the further translation that is needed in order to make the resulting PDDL2.1 formulation aware of the presence of events, for both cases.

Event-aware EXP Translation. Let $\Pi = \langle hF, X, I, G, A, E, P \rangle$ be a PDDL+ problem, and let $\Pi_{EXP} = \langle hF, X, I, G, A [fSIMg, ci$ the PDDL2.1 problem obtained using EXP ignoring set E , respectively. The handling of events E can be achieved by a further translation into $\Pi_{EXP}^{events} = \langle hF [F_E, X, I [fsim-evg, G \wedge sim-ev, A^{\theta} [fev-check(SIM)g [fSIMEVg, ci$ with $A^{\theta} = \bigcup_{a \in A} ev-check(a)$.

Event-aware POLY Translation. Let $\Pi = \langle hF, X, I, G, A, E, P \rangle$ be a PDDL+ problem, and let $\Pi_{POLY} = \langle hF [D [fpauseg, X [X^{cp}, I, G \wedge pause, A_c [A_p [fstart, endg, ci$ the PDDL2.1 problem obtained using POLY ignoring set E , respectively. The handling of events E can be achieved by a further translation into $\Pi_{POLY}^{events} = \langle hF [D [fpauseg [F_E, X [X^{cp}, I [fsim-evg, G \wedge pause \wedge sim-ev, A_c^{\theta} [A_p [fstart^{\theta}, ev-check(end)g [fSIMEVg, ci$ with $A_c^{\theta} = \bigcup_{a \in A_c} ev-check(a)$ and $pre(start^{\theta}) = pre(start) \wedge sim-ev$ and $eff(start^{\theta}) = eff(start)$.

Example 3.3 (EXP Translation with Events - Continuing on Example 3.1). *In the following we show how the problem obtained using EXP in Example 3.1, i.e., Π_{EXP} , can be extended to take also the events modelling cars' collisions into account, i.e.,*

$E = \bigcup_{\substack{car;car^0 \ 2cars \\ car \neq car^0}} f_{crash-fast_{car;car^0}, crash-slow_{car;car^0}}g$. For the sake of conciseness, event aliases are shortened, too. For example, the $crash-fast_{car1;car2}$ event is shortened to f_{1-2} . The PDDL2.1 task we get taking events into account is $\Pi_{EXP}^{events} = hF [F_E, X, I [f_{sim-ev}, G \wedge : sim-ev, A^0 [fev-check(SIM), SIMEVg, ci$ where:

$$\begin{aligned}
 F_E &= f_{fired_{f_{1-2}}, fired_{f_{2-1}}, fired_{s_{1-2}}, fired_{s_{2-1}}}g \\
 A^0 &= \bigcup_{car \ 2 f_{car1;car2}g} fev-check(accelerate_{car}), ev-check(decelerate_{car}), ev-check(overtake_{car}), \\
 &\quad ev-check(return_{car}), ev-check(start_{car}), ev-check(stop_{car})g.
 \end{aligned}$$

In the following we show how actions in A are modified by using function $ev-check()$ over cars:

$$\begin{aligned}
 ev-check(accelerate_{car}) &= heng-on_{car} \wedge ha_{car} < A_{max}i \wedge : sim-ev, fhincrease, a_{car}, 1i, sim-evgi \\
 ev-check(decelerate_{car}) &= heng-on_{car} \wedge ha_{car} > A_{min}i \wedge : sim-ev, fhdecrease, a_{car}, 1i, sim-evgi \\
 ev-check(overtake_{car}) &= heng-on_{car} \wedge hv_{car} > 0i \wedge : ot_{car} \wedge : sim-ev, fot_{car}, sim-evgi \\
 ev-check(return_{car}) &= heng-on_{car} \wedge hv_{car} > 0i \wedge ot_{car} \wedge : sim-ev, f: ot_{car}, sim-evgi \\
 ev-check(start_{car}) &= h: eng-on_{car} \wedge : sim-ev, feng-on_{car}, sim-evgi \\
 ev-check(stop_{car}) &= heng-on \wedge hv_{car} = 0i \wedge ha_{car} = 0i \wedge : sim-ev, f: eng-on_{car}, sim-evgi.
 \end{aligned}$$

The same for SIM :

$$\begin{aligned}
 pre(ev-check(SIM)) &=: sim-ev \\
 eff(ev-check(SIM)) &= fW_{fm1}g, W_{fm2}g, W_{fm1;m2}g, sim-evg.
 \end{aligned}$$

In both cases, the $ev-check()$ function is used to make sure that whenever an A or SIM is performed then event testing has to necessarily be performed by executing at least one $SIMEV$.

Finally, the operator $SIMEV$ is defined as follows:

$$\begin{aligned}
 pre(SIMEV) &= sim-ev \\
 eff(SIMEV) &= \overbrace{fW_{trig}^{f_{1-2}}, W_{trig}^{f_{2-1}}, W_{trig}^{s_{1-2}}, W_{trig}^{s_{2-1}}}g [\overbrace{fW_{red}^{f_{1-2}}, W_{red}^{f_{2-1}}, W_{red}^{s_{1-2}}, W_{red}^{s_{2-1}}}g [W_{sat} [W_{?}
 \end{aligned}$$

where the conditional effects are defined as follows:

$$\begin{aligned}
 W_{trig}^{f_{1-2}} &= hd_{car1} \quad d_{car2} < D_{T_i} \wedge hd_{car1} \quad d_{car2} \quad 0i \wedge ot_{car1} \wedge ot_{car2} \wedge : crashed \triangleright fcrashedg \\
 W_{trig}^{f_{2-1}} &= hd_{car2} \quad d_{car1} < D_{T_i} \wedge hd_{car2} \quad d_{car1} \quad 0i \wedge ot_{car2} \wedge ot_{car1} \wedge : crashed \triangleright fcrashedg \\
 W_{trig}^{s_{1-2}} &= hd_{car1} \quad d_{car2} < D_{T_i} \wedge hd_{car1} \quad d_{car2} \quad 0i \wedge : ot_{car1} \wedge : ot_{car2} \wedge : crashed \triangleright fcrashedg \\
 W_{trig}^{s_{2-1}} &= hd_{car2} \quad d_{car1} < D_{T_i} \wedge hd_{car2} \quad d_{car1} \quad 0i \wedge : ot_{car2} \wedge : ot_{car1} \wedge : crashed \triangleright fcrashedg \\
 W_{red}^{f_{1-2}} &= hd_{car1} \quad d_{car2} < D_{T_i} \wedge hd_{car1} \quad d_{car2} \quad 0i \wedge ot_{car1} \wedge ot_{car2} \wedge : crashed \triangleright ffired_{f_{1-2}}g \\
 W_{red}^{f_{2-1}} &= hd_{car2} \quad d_{car1} < D_{T_i} \wedge hd_{car2} \quad d_{car1} \quad 0i \wedge ot_{car2} \wedge ot_{car1} \wedge : crashed \triangleright ffired_{f_{2-1}}g \\
 W_{red}^{s_{1-2}} &= hd_{car1} \quad d_{car2} < D_{T_i} \wedge hd_{car1} \quad d_{car2} \quad 0i \wedge : ot_{car1} \wedge : ot_{car2} \wedge : crashed \triangleright ffired_{s_{1-2}}g \\
 W_{red}^{s_{2-1}} &= hd_{car2} \quad d_{car1} < D_{T_i} \wedge hd_{car2} \quad d_{car1} \quad 0i \wedge : ot_{car2} \wedge : ot_{car1} \wedge : crashed \triangleright ffired_{s_{2-1}}g \\
 W_{satu} &= (pre(f_{1-2}) _ fired_{f_{1-2}}) \wedge (pre(f_{2-1}) _ fired_{f_{2-1}}) \wedge (pre(s_{1-2}) _ fired_{s_{1-2}}) \wedge \\
 &\quad (pre(s_{2-1}) _ fired_{s_{2-1}}) \triangleright \mathcal{F}: sim-ev, : fired_{f_{1-2}}, : fired_{f_{2-1}}, : fired_{s_{1-2}}, : fired_{s_{2-1}}g.
 \end{aligned}$$

3.4 Properties

In the following theorem, we prove the soundness and completeness for POLY and EXP translations in the general case for PDDL+ problems with events by using Lemmas 1 and 2.

Theorem 1 (Soundness and Completeness of POLY and EXP). *Let $\Pi = \langle hF, X, I, G, A, E, P \rangle$ be a PDDL+ planning instance, and let Π_{EXP}^{events} (Π_{POLY}^{events}) be the PDDL2.1 planning instance obtained by using the EXP (POLY) translation. Π admits a solution under δ discretisation iff so does Π_{EXP}^{events} (Π_{POLY}^{events}).*

Proof. We focus on Π_{EXP}^{events} (the proof for Π_{POLY}^{events} is similar) and prove the two directions largely exploiting the constructions devised for the event-free translations.

(\Leftarrow) Let $ES(T_i) = \langle hE_{trigg}(T_i), \dots, E_{trigg}(T_{i+k-1}) \rangle$ be the unique and finite⁷ cascade of events from an STP T_i of the discrete projection H of π_t solving Π .

We show that it is possible to construct a plan π_{EXP} from π_t such that π_{EXP} is valid for Π_{EXP} . Differently from an event-free planning task where we can define a mapping from π_t to π_{EXP} , in this case we have to resort to H to build a valid π_{EXP} . Let $T_H = \langle hT_0, \dots, T_m \rangle$ be the $m+1$ STPs of H , we define π_{EXP} in two steps. As a first step, we define the sequence $\pi_{EXP}^0 = \langle ha_0^0, \dots, a_{m-1}^0 \rangle$ such that for all $i \in [0..m-1]$

$$a_i^0 = \begin{cases} a_i & \text{if } H_A(T_i) = ha_j \notin hi \\ SIMEV & \text{if } E_{trigg}(T_i) \notin hi \\ SIM & \text{otherwise} \end{cases}$$

Then we obtain π_{EXP} from π_{EXP}^0 by inserting a *SIMEV* action just before any action $a \in \pi_{EXP}^0$ such that $a \notin SIMEV$, and just before the end of the plan. Let $\tau^0 = \langle hs_0, \dots, s_m \rangle$

7. Note that, under the restriction imposed over PDDL+, for each $E; E^0 \in set(ES)$ with $E \notin E^0$, $set(E) \setminus set(E^0) = \{ \}$; and, for each $''; ''^0 \in set(E)$ with $E \in set(ES)$, $''$ and $''^0$ are not mutex, as long as $'' \notin ''^0$.

be the sequence of states obtained by applying iteratively actions from π_{EXP} and filtering out those states produced by any last *SIMEV* of a series, the difficult bit is to show that τ^θ is equivalent to $\tau = hH_S(T_0), \dots, H_S(T_m)i$ under variables $F [X$. The first observation is that $j\tau j = j\tau^\theta j$; and this follows directly from the fact that the number of states we are filtering out is exactly the number of *SIMEV* that we have added to the plan. Then, in order to prove that $H_S(T_i)$ and s_i are equivalent for all $i \geq [0..m]$, we can use the same arguments of Lemma 1 extended to account for the case where the transition is due to a cascade of events. More precisely, let $T_i = ht_i, n_i i$ be a STP in which a cascade of events $ES(T_i)$ with $jES(T_i)j = k$ is triggered. By using R1 of Definitions 13-16 we know that, for each $j \geq [i..i+k]$ $H_S(T_j) \not\equiv \bigwedge_{e \in E_{\text{trigg}}(T_j)} \text{pre}(e)$, $H_S(T_{j+1}) = \gamma(H_S(T_j), E_{\text{trigg}}(T_j))$ and $T_{j+1} = ht_i, n_j + 1 i$.

In order to prove $H_S(T_{i+k}) = s_{i+k}$ with $s_{i+k} = \gamma(s_i, \text{SIMEV } k)$ we need to show that, for every j such that $j \geq [i..i+k]$, $H_S(T_j) = s_j$; we do so, again, by induction. The base case ($j = i$) is trivially proved (inherited by Lemma 2). For the inductive step it suffices to observe that *SIMEV* exhibits a behaviour that is equivalent to $h>$, $\bigcup_{\substack{c.e2e \text{ ("} \\ \text{with "2Etrigg}(T_j)}}$ $c \triangleright ei$ and

then it follows that $H_S(T_{j+1})$ and $s_{j+1} = \gamma(s_j, \text{SIMEV})$ are equivalent. The only thing that is missing is to show that when each action is applied in π_{EXP} the variable *sim-ev* is false. But this directly follows from the fact that the last *SIMEV* is applied when all events have been triggered. Indeed, we have that $E_{\text{trigg}}(T_{i+k}) = hi$, and $H_S(T_{i+k}) = s_{i+k}$. So *SIMEV* will make *sim-ev* false and get ready for the next round of execution by resetting all the monitoring variables *fired* to false.

(()) The mapping from π_{EXP} to π_t is identical to what was done for Lemma 1, but for the fact that all occurrences of *SIMEV* are ignored. Then, the π_t plan can be proved valid against the discretised PDDL+ model by observing that, for each series of *SIMEV* of length k the projection H encompasses $k - 1$ STPs, one for each *SIMEV* that triggers a change on variables in $F [X$. Each STP generates a set of events whose effects are those that arise from the active conditional effects of the associated *SIMEV*. This is due to the fact that the *SIMEV* conditional effect's condition subsumes the precondition of each event associated with it. □

In the following theorems, we show how the translations EXP and POLY impact the size of the compiled problems Π_{EXP} and Π_{POLY} .

Theorem 2 (Size of Π_{EXP}). *Let $\Pi = hF, X, I, G, A, ;, Pi$ be a PDDL+ planning instance the reformulation POLY produces a PDDL2.1 planning instance $\Pi_{\text{EXP}} = hF, X, I, G, A [fSIMg, ci$ that increases the size of Π exponentially.*

Proof. In the compiled problem Π_{EXP} the set of actions is extended with the *SIM* action which has a number of conditional effects equal to $jP^+(P)j = jP(P) n : j = 2^P - 1$. □

Theorem 3 (Size of Π_{POLY}). *Let $\Pi = hF, X, I, G, A, ;, Pi$ be a PDDL+ planning instance the reformulation EXP produces a PDDL2.1 planning instance $\Pi_{\text{POLY}} = hF [D [fpauseg, X [X^{cp}, I, G [f : pauseg, A_c [A_p [fstart, endg, ci$ that increases the size of Π only polynomially.*

Proof. The propositional variables increase by $|D| + 1$ where, in the worst case, $|D| = |P| + |X|$. Numeric variables are doubled; indeed $|X^{cp}| = |X|$. Finally, since $|A_c| = |A|$, the actions increase by $|A_p| + 2$ therefore, in the worst case, $|A_p| = |P| + |X|$. \square

In his study, Nebel (2000) considered the effects of the translations between planning formalisms on the size of the plans solving the reformulated instances, besides the needed temporal and spatial resources.

In our context, given a PDDL+ problem Π and a reformulation $Z \in \{POLY, EXP\}$, we say that Z *preserves the plan size exactly*, up to additive constants, iff for each plan $\pi_t = \langle h, \pi, ht_s, t_e \rangle$ which solves Π there exists a PDDL2.1 plan π_Z such that $|\pi_Z| = |\pi| + k$ where $k \in \mathbb{N}_0$; Z *preserves the plan linearly* iff for each plan $\pi_t = \langle h, \pi, ht_s, t_e \rangle$ which solves Π there exists a PDDL2.1 plan π_Z such that $|\pi_Z| = c|\pi| + k$ where $c, k \in \mathbb{N}_0$; finally, Z *preserves the plan polynomially* iff for each plan $\pi_t = \langle h, \pi, ht_s, t_e \rangle$ which solves Π there exists a PDDL2.1 plan π_Z such that $|\pi_Z| = p(|\Pi|, |\pi|)$ where $p(\cdot)$ is a polynomial expression that depends on the size of Π , i.e., $|\Pi|$, and on the size of π , i.e., $|\pi|$.

Theorem 4 (Size of Plans for Π_{EXP} and Π_{POLY}). *Let $\Pi = \langle h, F, X, I, G, A, \delta, P \rangle$ be an event-free PDDL+ planning instance, and let Π_{POLY} and Π_{EXP} be the PDDL2.1 obtained by using the POLY and EXP translations, respectively. The reformulations POLY and EXP preserve the plan size polynomially and linearly, respectively.*

Proof. Let $\pi_t = \langle h, \pi, ht_s, t_e \rangle$ be a solution for Π under δ discretisation and let π_{POLY} and π_{EXP} be the corresponding plan for Π_{POLY} and Π_{EXP} , respectively. Using the rules provided in Lemmas 1-2 for mapping π_t into π_{EXP} and π_{POLY} respectively, we get that for each discrete advance of time in the original plan, i.e., $\frac{t_e - t_s}{\delta}$, we have to execute the action *SIM* in π_{EXP} and the sequence of actions *wait* in π_{POLY} . Then, for EXP we obtain:

$$|\pi_{EXP}| = |\pi| + \frac{t_e - t_s}{\delta}$$

Each *wait* in π_{POLY} consists of two delimiting actions, i.e., *start* and *end*, plus an action for each numeric effect of each process; in the worst case, each process has a numeric effect for each variable of the problem. By adding these contributions we obtain:

$$|\pi_{POLY}| = |\pi| + \frac{t_e - t_s}{\delta} + (|P| + |X| + 2)$$

\square

Theorem 5 (Size of Plans for Π_{EXP}^{events} and Π_{POLY}^{events}). *Let $\Pi = \langle h, F, X, I, G, A, E, P \rangle$ be a PDDL+ planning instance, and let Π_{EXP}^{events} and Π_{POLY}^{events} be the PDDL2.1 obtained by using the POLY and EXP translations, respectively. The reformulation POLY and EXP preserve plan size polynomially.*

Proof. Let $\pi_t = \langle h, \pi, ht_s, t_e \rangle$ be a solution for Π under δ discretisation and let π_{POLY} and π_{EXP} be the corresponding plan for Π_{POLY}^{events} and Π_{EXP}^{events} , respectively.

Using the rules outlined in Lemma 2 and Theorem 1 to map π_t into π_{POLY} we can prove the upper-bound on $|\pi_{POLY}|$ as follows: (i) each action of π_{POLY} has to be followed by at least one *SIMEV* action up to a maximum of $|E| + 1$ and such sequence has to be executed also at

the beginning of π_{POLY} since $I^\theta \not\models \text{sim-ev}$; (ii) possible event triggers must be checked after each block of actions simulating the passage of time, i.e., the *wait* sequence; therefore the term $jEj+1$ has to be multiplied by the number of time steps occurred within the envelope ht_s, t_ej , i.e., $\frac{t_e}{\delta} t_s$; (iii) each *wait* in π_{POLY} consists of two delimiting actions, i.e., *start* and *end*, plus an action for each numeric effect of each process; in the worst case, each process has a numeric effect for each variable of the problem. By adding these contributions we obtain:

$$j\pi_{\text{POLY}}j \quad (j\pi j+1) \quad \overbrace{(jEj+1)}^{(i)} + \frac{t_e}{\delta} t_s \quad \overbrace{((jEj+1) + (jPj \ jXj+2))}^{(ii)} \quad \overbrace{1}^{(iii)}$$

The only difference for π_{EXP} is that instead of the *wait* sequence for the contribute iii) we apply the single action *SIM*; then, we obtain:

$$j\pi_{\text{EXP}}j \quad (j\pi j+1) \quad \overbrace{(jEj+1)}^{(i)} + \frac{t_e}{\delta} t_s \quad \overbrace{((jEj+1) + 1)}^{(ii)} \quad \overbrace{1}^{(iii)}$$

□

4. Optimising POLY and EXP

This section presents two optimisations for our translation schemata, both aimed at reducing the length of valid plans needed to solve the translated problem. Our optimisations preserve both the soundness and the completeness of the approach; the basic idea is to exploit conditions obtained by looking at the structure of our problem. Such conditions intercept whether some action or an event does not trigger any other event. Our first optimisation allows us to prune the *SIMEV* action in case an action cannot be followed by some event. Thanks to this optimisation our translated problem can potentially chain a sequence of actions before synchronising their effects with exogenous events. The second optimisation is aimed at simplifying the *SIMEV* definition through an analysis of the structure of all events. Both optimisations are constructed by introducing the notion of Trigger-Free action (or event), which is under-approximated, by intercepting a sufficient condition obtained through reasoning by regression.

4.1 Avoiding Events Checking

In Section 3.3 we have seen that each action is modified in order to ensure that after its execution, a *SIMEV* sequence is forced to be executed. This is obtained by using the function *ev-check*() that takes as input an action a , and returns another action a^θ where $pre(a^\theta) = pre(a) \wedge : \text{sim-ev}$ and $eff(a^\theta) = eff(a) [\text{fsim-evg}$.

The rationale behind the first optimisation we present consists of studying conditions under which an action does not need to be followed by the potentially expensive *SIMEV* sequence. To do so, we introduce the notions of *Trigger-Free* and *Universally Trigger-Free* actions.

Definition 18 (Trigger-Free and Universally Trigger-Free action). *Given a PDDL+ problem Π , let $a \in A$ and let $\varepsilon \in E$, we say that a is Trigger-Free w.r.t. ε , denoted with $TF(a, \varepsilon)$, iff for each state s such that for all $\varepsilon^\theta \in E$, $s \not\models pre(\varepsilon^\theta)$, then $\gamma(s, a) \not\models : pre(\varepsilon)$. Moreover,*

we say that a is Universally Trigger-Free, denoted with $UTF(a)$, iff a is Trigger-Free w.r.t. all events in E , i.e., $UTF(a), \forall \varepsilon \in E, TF(a, \varepsilon)$.

Intuitively, this notion captures all those actions that do not change the state in a way that some event can be triggered. The UTF definition can be used to slightly modify either POLY or EXP to disallow the application of function $ev-check()$ to all UTF actions. For example, for the event-aware variant of EXP, the definition of $A^0 = \bigcup_{a \in A} ev-check(a)$ is replaced with:

$$A^0 = \{ ev-check(a) \mid a \in A \wedge UTF(a) \} \cup \{ a \mid a \in A, UTF(a) \}$$

The optimised variant of the event-aware variant of POLY is straightforward.

It is worth noting that, considering our very expressive language for modelling actions, the computation of an exact Trigger-Free relation is complicated. Firstly, because we have numeric effects and propositional effects. Secondly, because any such effect can depend on the state in which the action is applied. Thirdly, because we allow the combination of any NNF formula in the precondition of the events. For these reasons, in what follows we propose an under-approximation of such a relation.

We under-approximate the Trigger-Free relationship through Algorithm 1, which works by sequencing two checks. The first check (Line 2) is a neutrality test that evaluates whether a does not affect anything that is involved in the preconditions of ε . If that is the case, then there is no way ε can be triggered. This check is done by inspecting if the action affects any variable in some necessary condition for the precondition of ε ; formally $vars(eff(a)) \setminus vars(necessary(pre(\varepsilon))) \neq \emptyset$, where:

- given a precondition expressed as a formula φ , the function $necessary(\varphi)$ returns those terms that have to necessarily be true in order to ensure that φ is true. We restrict the attention to top-level conjuncts. For example, consider the case in which $\varphi = p \wedge :q \wedge hx > 10i \wedge (c _ (b \wedge hy + z < 10i)$, the necessary conditions are that p and $:q$ hold true, and x is greater than 10, i.e., $necessary(\varphi) = \{p, :q, hx > 10i\}$;
- $vars()$ is a function that returns the set of all variables, whether propositional or numeric, involved in its parameter; to be specific, if the parameter is a formula φ then $vars(\varphi)$ returns all the variables exploring recursively the formula, e.g., given $\varphi = p \wedge hx > 10i$, then $vars(\varphi) = \{p, x, g\}$; if the parameter is an effect e then $vars(e)$ returns all the affected variables, e.g., if $e = f > \triangleright fpg, > \triangleright fhinc, x, 20 + yigg$, then $vars(e) = \{p, x, g\}$.

The second check is performed by using a sufficient condition that checks whether action a always makes the preconditions of ε unsatisfied (Lines 4-6). The procedure seeks if there exists a necessary conjunct g of $pre(\varepsilon)$ such that the logical conjunction of the formula resulting by applying operator R (see below) on g and some necessary preconditions of the action generates a contradiction (Line 5). For example, if there is an operator R that results in $hx > 0i$, and the action has one precondition which states that $hx > 0i$, then the conjunction $hx > 0i \wedge hx < 0i$ is not satisfiable and therefore the procedure returns TRUE.⁸. The operator R is similar to a regression function; it takes as an input an action and a

8. For this check we rely on an external solver; in our experiment, we used SIMPY (Meurer et al., 2017).

condition (propositional or numeric condition) and returns a formula, i.e., $R : A \rightarrow \Lambda \mid \neg \Lambda$, where Λ is the set of conditions (propositional and numeric) that can be expressed in PDDL+. It is formally defined as follows:

$$R(a, g) = \begin{cases} g & \text{if } \text{vars}(g) \setminus \text{vars}(\text{eff}(a)) = \emptyset \\ > & \text{if } g = hf = bi \text{ and } \exists c \triangleright f\dots, hf := bi \text{ g } \not\geq \text{eff}(a) \\ ? & \text{if } g = hf = bi \text{ and } \exists c \triangleright f\dots, hf := b^0 i \text{ g } \geq \text{eff}(a) \text{ and } b \not\leq b^0 \\ NR(a, g) & \text{if } g = h\xi \bowtie 0i \text{ and } \text{direct}(a, g) \\ > & \text{if } g = h\xi \bowtie 0i \text{ and not } \text{direct}(a, g) \end{cases} \quad (5)$$

In the formula, the function *direct* returns true for all actions whose numeric effects on g are unconditionally triggered, i.e., the left-hand side of all numeric conditional effects onto g is $>$. Formally: $\text{direct}(a, g) = \exists c \triangleright e \geq \text{eff}(a) \wedge g \geq e, c = >$.

NR is the effect regressor as for Scala, Haslum, Thiébaux, and Ramírez (2020), slightly reformulated to consider conditional effects as an input.⁹ More precisely, NR transforms any numeric condition $\xi \bowtie 0$ in $\xi[x_1/\tau(a, x_1), \dots, x_k/\tau(a, x_k)] \bowtie 0$ where

$$\tau(a, x_i) = \begin{cases} \xi^0 & \text{if } \exists c \triangleright f\dots, \text{hasgn}, x_i, \xi^0 i \text{ g } \geq \text{eff}(a) \\ x_i + \xi^0 & \text{if } \exists c \triangleright f\dots, \text{hinc}, x_i, \xi^0 i \text{ g } \geq \text{eff}(a) \\ x_i - \xi^0 & \text{if } \exists c \triangleright f\dots, \text{hdec}, x_i, \xi^0 i \text{ g } \geq \text{eff}(a) \\ x_i & \text{Otherwise} \end{cases} \quad (6)$$

and x_1, \dots, x_k are the k numeric variables affected by a . The $/$ operator denotes a substitution operator for manipulating numeric conditions, e.g., for $ha + b > 0i$ then $\xi[a/(a + 1), b/(b + 1)] = h(a + 1) + (b + 1) > 0i$.

Intuitively, Equation 5 distinguishes whether the given input is a propositional condition or a numeric one, and regresses failure (i.e., $?$) when the action makes unsatisfied the condition independently on the state in which the action is applied. In order to be sure that the action *always* makes a condition violated (for instance deletes it), operator R conservatively excludes all cases where the action *may* achieve the condition. In this latter case, the function safely returns $>$. If the action is however not affecting the condition in any sensible manner, i.e., the first case of the equation, the condition is left unaltered.

Lemma 3. *Let a be an action and ε be an event. If Algorithm 1 returns TRUE the relation $TF(a, \varepsilon)$ holds.*

Proof. Algorithm 1 returns TRUE if the condition at Line 2 holds or if at least one necessary condition within the precondition of ε will not hold after the execution a . If the action does not interact with any of the variables in the precondition of ε , then the action is Trigger-Free w.r.t. ε in that if ε is not triggered in some state, it is certainly not a that will make it active. Regarding the second condition, it suffices to observe that if a makes certainly false at least one of its necessary preconditions, or modifies them in a way that when used

9. Note that conditional effects here are considered only syntactically, but are basically ignored. How to extend regression with conditional effects with numeric effects in order to handle them in a sensible manner is out of the scope of this work.

Algorithm 1: Algorithm for under approximating when a is Trigger-Free w.r.t. an event ε

Input: an action or an event a and an event ε
Output: a Boolean value

```

1 Function  $TF(a, \varepsilon)$ :
2   if  $\text{vars}(\text{eff}(a)) \setminus \text{vars}(\text{necessary}(\text{pre}(\varepsilon))) = \emptyset$  then
3     return TRUE
4   for  $g \in \text{necessary}(\text{pre}(\varepsilon))$  do
5     if  $R(a, g) \wedge \bigwedge_{g' \in \text{necessary}(\text{pre}(a))} g' \neq ?$  then
6       return TRUE
7   end
8   return FALSE
    
```

in conjunction with the precondition of a the arising satisfiability problem is unsatisfiable, then there is no way event ε could be triggered after executing a . Therefore, also in this case, returning TRUE implies that action a is Trigger-Free w.r.t. ε . \square

Theorem 6. *The optimisation variants of EXP and POLY using Algorithm 1 to under-approximate the Trigger-Free relation preserves the soundness and the completeness of both translations.*

Proof. The proof for this theorem follows directly from Definition 18. Indeed, we are basically only avoiding the execution of events that can never be applied after a Trigger-Free action, and moreover, by Lemma 3 we know that Algorithm 1 only returns TRUE if the Trigger-Free relation holds. \square

In the following example, we detail how Trigger-Free relations are evaluated in practice.

Example 4.1 (Trigger-free Operators). *Let $\varepsilon = hx = 10i \wedge hy + z > 20i \wedge hw = >i, >\triangleright fha := ?igi$ be an event and let a_1, a_2 and a_3 be three actions such that:*

$$\begin{aligned}
 a_1 &= h>, f>\triangleright fha := >ig, >\triangleright fhinc, b, 10igg \\
 a_2 &= h>, f>\triangleright fhw := ?igg \\
 a_3 &= hhy < 10i, f>\triangleright fhinc, y, 10ig, >\triangleright fhasgn, z, 0igg
 \end{aligned}$$

We study if the property $TF(a, \varepsilon)$ holds for each $a \in \{a_1, a_2, a_3\}$ by using step by step Algorithm 1.

To establish that a_1 is Trigger-Free w.r.t. ε , it suffices the first check of Algorithm 1 (Line 2). Since $\text{vars}(\text{eff}(a_1)) = \text{vars}(f>\triangleright fha := >ig, >\triangleright fhinc, b, 10igg) = \{a, b, g\}$ and $\text{vars}(\text{necessary}(\text{pre}(\varepsilon))) = \text{vars}(\text{pre}(\varepsilon)) = \text{vars}(hx = 10i \wedge hy + z > 20i \wedge hw = >i) = \{x, y, w, z, g\}$, then $TF(a_1, \varepsilon)$ holds. Indeed, the action is neutral w.r.t. the event.

Concerning a_2 , since $\text{vars}(\text{eff}(a_2)) = \{w, g\}$, then $\text{vars}(\text{eff}(a_2)) \setminus \text{vars}(\text{pre}(\varepsilon)) = \{w, g\} \setminus \{x, y, w, z, g\} = \{w, g\} \neq \emptyset$; and therefore a_2 is not neutral w.r.t. ε . It is necessary to investigate further what kind of relationship between a_2 and ε exists. Let $\text{pre}(\varepsilon) = g_1 \wedge g_2 \wedge g_3$ where $g_1 = hx = 10i$, $g_2 = hy + z > 20i$ and $g_3 = hw := ?i$, then $\text{necessary}(\text{pre}(\varepsilon)) = \{g_1, g_2, g_3\}$.

Before proceeding, note that for a_2 we have $\bigwedge_{g^0 \in \text{necessary}(\text{pre}(a_2))} g^0 = \text{pre}(a_2) = >$. We iterate over the necessary conditions $f g_1, g_2, g_3 g$:

- $g_1 = hx = 10i$; since $\text{vars}(g_1) = fxg$ and $\text{vars}(\text{eff}(a_2)) = fwg$ then $\text{vars}(g_1) \setminus \text{vars}(\text{eff}(a_2)) = ;$; therefore we fall into case (i) of Formula 5. Therefore $R(a_1, g_1) \wedge \text{pre}(a_2) = g_1 \wedge > = hx = 10i \wedge > \not\equiv ?$;
- $g_2 = hy + z > 20i$; analogously to what seen for g_1 , we get that: $R(a_2, g_2) \wedge \text{pre}(a_2) = hy + z > 20i \wedge > \not\equiv ?$;
- $g_3 = hw = >i$; since g_3 is a propositional condition involving w and there exists a propositional assignment $> \triangleright fhw := ?ig \not\equiv \text{eff}(a_2)$ such that the two Boolean values are in opposition, we fall into the case (iii) of Formula 5. Therefore $R(a_2, g_3) \wedge \text{necessary}(\text{pre}(a_2)) = ? \wedge > \not\equiv ?$;

We have shown that there is at least one necessary condition of $\text{pre}(\varepsilon)$, i.e., g_3 , which is always falsified by a_2 , therefore we have proved that $TF(a_2, \varepsilon)$ holds.

As far as it is concerned by a_3 , since $\text{vars}(\text{eff}(a_3)) = \text{vars}(f > \triangleright fhinc, y, 10ig, > \triangleright fhasgn, z, 0igg) = fy, zg$, then $\text{vars}(\text{eff}(a_3)) \setminus \text{vars}(\text{pre}(\varepsilon)) = fy, zg \not\equiv ;$ and therefore a_3 is not neutral with respect to ε . Before proceeding, note that for a_3 we have that

$\bigwedge_{g^0 \in \text{necessary}(\text{pre}(a_3))} g^0 = \text{pre}(a_3) = hy < 10i$. Again, we iterate over the necessary conditions of $\text{pre}(\varepsilon)$:

- $g_1 = hx = 10i$; for g_1 , following the same steps above for a_2 , we obtain $R(a_3, g_1) \wedge \text{pre}(a_3) = hx = 10i \wedge hy < 10i \not\equiv ?$;
- $g_2 = hy + z > 20i$; since $\text{vars}(g_2) \setminus \text{vars}(\text{eff}(a_3)) \not\equiv ;$, $\text{direct}(a_3, g_2)$ holds and g_2 is a numeric condition, we fall into the case (iv) of Formula 5; we get that $R(a_3, g_2) = NR(a_3, hy+z > 0i) = h(y+z - 20)[y/y+10, z/0] > 0i = hy+10 - 20 > 0i = hy - 10 > 0i$; finally, we get that $R(a_3, g_2) \wedge \text{pre}(a_3) = hy > 10i \wedge hy < 10i \not\equiv ?$;

We can stop the iteration, ignoring the evaluation of g_3 , as we have shown that there is a necessary condition of $\text{pre}(\varepsilon)$ which is always falsified by the execution of a_3 and therefore $TF(a_3, \varepsilon)$ holds.

4.2 Avoiding Cascade of Events Handling

Another optimisation can also be applied to the *SIMEV* action. The idea is to detect if a set of events can not yield a cascade of events. We achieve this by slightly reinterpreting the Trigger-Free relation among events instead of that between an action and an event. That is:

Definition 19 (Trigger-Free and Universally Trigger-Free Event). *Given a PDDL+ problem, let $\varepsilon, \varepsilon^0 \in E$ and let $\varepsilon \in E$, we say that ε is Trigger-Free w.r.t. ε^0 , namely $TF(\varepsilon, \varepsilon^0)$ iff for each state s where all $\varepsilon^0 \in E$ are such that $s \not\equiv \text{pre}(\varepsilon^0)$, we have that $\gamma(s, \varepsilon) \not\equiv : \text{pre}(\varepsilon^0)$. Moreover, we say that ε is Universally Trigger-Free, denoted with $UTF(\varepsilon)$, iff ε is Trigger-Free w.r.t. all events in E , i.e., $UTF(\varepsilon)$, $\forall \varepsilon^0 \in E, TF(\varepsilon, \varepsilon^0)$.*

Note that under our assumptions, $TF(\varepsilon, \varepsilon)$ always holds since events have to self-deactivate.

We use the Trigger-Free notion to prevent the application of the expensive conditional effects W_{red} and W_{satu} within $SIMEV$ used for handling the cascade of events. More precisely, these conditional effects can be avoided if for each $\varepsilon \in E$ then $UTF(\varepsilon)$ holds.

Therefore $SIMEV$ can be formulated by adding a switch that controls the relationship between each pair of events:

$$\begin{aligned} pre(SIMEV) &= sim-ev \\ eff(SIMEV) &= \begin{cases} W_{trig} [W_{\text{?}} [f > \triangleright f: sim-evgg & \text{if } \exists \varepsilon \in E, UTF(\varepsilon) \text{ holds} \\ W_{trig} [W_{red} [W_{satu} [W_{\text{?}} & \text{otherwise.} \end{cases} \end{aligned}$$

The Trigger-Free relation can be under-approximated using Algorithm 1, pretty much verbatim.

Theorem 7. *The optimisation of EXP and POLY using the generalised SIMEV preserves the soundness and the completeness of both of the translations.*

Proof. It suffices to observe that W_{red} and W_{satu} only serve the purpose of tracking which event has been already triggered and whether all events have been triggered. As none of the events triggers any other event, this is not necessary, so we can stop after exactly one execution of $SIMEV$. \square

4.3 Translations Notation

In this section, we clarify the notation used in the translation presented in this work and its relationship with the translations presented in our previous work (Percassi et al., 2021). Given a translation $Z \in \{EXP, POLY\}$, we denote with:

- Z_0 the plain translation without any optimisation;
- Z_1 the translation using the optimisation described in Section 4.2;
- Z_2 the translation using the optimisation described in Section 4.1;
- Z_3 the translation jointly using the optimisations described in Sections 4.1 and 4.2.

In our previous work (Percassi et al., 2021), EXP and POLY used in the experiments correspond to EXP_1 and $POLY_1$.

5. Experimental Analysis

Our experimental analysis aims at assessing the extent to which the introduced translations allow to reformulate PDDL+ instances into instances amenable for PDDL2.1 planning engines, and investigating the impact of the described optimisations.

5.1 Experimental Settings

We consider three engines at the state of the art for PDDL+ planning: ENHSP version 20 (Scala et al., 2020) with the additive interval-based relaxation heuristic (Scala, Haslum, Thiébaux, & Ramírez, 2016), SMTPLAN (Cashmore et al., 2020), DiNO (Piotrowski, Fox, Long, Magazzeni, & Mercurio, 2016) and UPMURPHI (Penna, Magazzeni, & Mercurio, 2012). As a PDDL2.1 planning engine we use the well-known METRIC-FF (Hoffmann, 2003). None of the considered systems requires the user to provide a bound on the length of valid plans. We did not consider other numeric planning systems such as LPG (Gerevini, Saetti, & Serina, 2008), OPTIC (Benton, Coles, & Coles, 2012) or the same ENHSP because none of them provides effective support for conditional effects and negative preconditions. All the planning engines have been run using default parameters. Numeric planners are used on numeric instances obtained using the proposed translations with $\delta = 1$. PDDL+ planners which reason over a discrete timeline are used with $\delta = 1$. Finally, to study how the translations affect the makespan w.r.t. a reference, we compared the makespan of the plans found with ENHSP used with the A (h^{blind}) heuristic and those found with METRIC-FF on all numeric tasks.

Our experiments were run on an Intel Xeon Gold 6140M CPU with 2.30 GHz. For each instance, we set a cutoff time of 900 seconds, and memory was limited to 8 GB.

For our experimental evaluation, we consider six benchmark domains. Three of them, LINEAR-CAR (LIN-CAR), LINEAR-GENERATOR (LIN-GEN), and SOLAR-ROVER (ROVER), are well-known PDDL+ benchmarks. OVERTAKING-CAR (OT-CAR) is a version of LINEAR-CAR that extends the original domain by considering multiple lanes, and the need for the car to move between lanes in order to avoid obstacles (see Example 2.3 for a description of the model of this domain). BAXTER (Bertolucci, Capitanelli, Maratea, Mastrogiovanni, & Vallati, 2019) and URBAN-TRAFFIC-CONTROL (UTC) (Vallati, Magazzeni, Schutter, Chrapa, & McCluskey, 2016; McCluskey & Vallati, 2017) are taken from real-world applications. The BAXTER domain exploits planning for supporting robots in dealing with articulated object manipulation tasks. The UTC domain models the use of planning for generating traffic light signal plans in order to de-congest an area of an urban region.

Our implementation of the translator is written in PYTHON 3 and makes use of the SYMPY library (Meurer et al., 2017) for solving the system of equations that arises for establishing Trigger-Free actions. The benchmark suite and the tool for translating PDDL+ instances are available at <https://bit.ly/30gMyNW>.

5.2 Size of the Translated Instances

First, we turn our attention to the size increase that can result from the use of the proposed translations. A direct comparison is not possible, as the original models and the translated models are encoded using different languages. For this reason, we introduce a notion of size increase ratio as follows. Let $\Pi = \langle hF, X, I, G, A, P, E \rangle$ be a PDDL+ problem and let $\Pi_T = \langle hF^\theta, X^\theta, I^\theta, G^\theta, A^\theta, c \rangle$ the corresponding PDDL2.1 problem obtained by using translation T . If $\hat{\Pi}$ is a PDDL2.1 problem, we define the size increase ratio, denoted with r , introduced by T as:

$$r(T) = \frac{jA^\theta j + jW^\theta j}{jA j + jP j + jE j + jW j}$$

Domain	(P) (E)		POLY				EXP			
			POLY ₃	POLY ₂	POLY ₁	POLY ₀	EXP ₃	EXP ₂	EXP ₁	EXP ₀
ROVER (20)	4.0	5.0	1.9	1.9	1.9	1.9	1.8	1.8	1.8	1.8
LIN-CAR (10)	2.0	0.0	1.7	1.7	1.7	1.7	1.3	1.3	1.3	1.3
LIN-GEN (10)	6.1	8.3	2.0	2.0	2.5	2.5	16.0	16.0	16.5	16.5
UTC (10)	34.1	15.8	2.5	2.5	2.8	2.8	16384.4	16384.4	16384.8	16384.8
BAXTER (10)	56.0	22.0	1.5	1.5	1.7	1.7	—	—	—	—
OT-Car (20)	4.1	5.4	1.3	1.3	1.6	1.6	2.1	2.1	2.4	2.4

Table 1: For each domain, $\mu(jPj)$ and $\mu(jEj)$ denote the average number of grounded processes and events, respectively, while $r(\text{POLY})$ and $r(\text{EXP})$ denote the average size increase ratio of the instances. Between brackets, the number of problem instances considered for each benchmark domain. Symbol “—” indicates a translation failure due to memory limits. “ ” indicates that in the UTC domain we only consider 4 instances out of 10; the translation failed due to memory limits in the remaining instances.

where \mathcal{W} and \mathcal{W}^0 denote the set of conditional effects of Π and Π_T , respectively. In other words, we measure the size of the PDDL+ model in terms of actions, processes, and events, and we measure the size of the corresponding PDDL2.1 model in terms of actions and the number of conditional effects introduced by the compilation. This is because conditional effects play a major role in the proposed translations, and can be challenging to deal with by planners.

Table 1 provides an overview of the average number of processes and events of the considered benchmark domains, and the increased ratio obtained by using the proposed translations with/without the optimisations. The EXP translations are significantly larger than the POLY translations, regardless of the optimisation. This is not the case in domains where the number of processes and events is very limited, such as LINEAR-CAR and SOLAR-ROVER, where the EXP models are smaller than the POLY ones. Intuitively, this is due to the fact that when the number of processes is very small, the additional actions needed by the POLY translation can lead to a larger model. The use of such actions is instead beneficial in large instances, where the EXP approach can sometimes blow up the memory budget.

Considering the size increase ratio, the POLY₂ optimisation can usually lead to smaller instances for both the considered translations.

5.3 Assessment of Performance

Table 2 shows the performance achieved by our translations, and optimisation on the considered benchmark domains when run using the METRIC-FF planning system. Results are shown in terms of coverage (number of solved instances), CPU-Time, quality of the generated plans (makespan), and number of nodes evaluates during the planning process. CPU-Time, quality, and evaluated nodes are presented as average over the instances solved by all the encodings under evaluation. Drawing a parallel with Table 1, it is easy to notice that in domains where there is a large number of processes and events, POLY translation delivers the best performance. The EXP translations are extremely large and hard to be dealt with by the planning engine. Conversely, the EXP translation seems to be more suitable for compact problems where the number of processes and events is limited. It is worth

Domain	Coverage							
	POLY				EXP			
	POLY ₀	POLY ₁	POLY ₂	POLY ₃	EXP ₀	EXP ₁	EXP ₂	EXP ₃
ROVER (20)	20	19	20	19	20	20	20	20
LIN-CAR (10)	10	10	10	10	10	10	10	10
LIN-GEN (10)	6	10	5	10	3	3	5	5
UTC (10)	0	7	0	7	0	0	0	0
BAXTER (20)	1	19	1	17	0	0	0	0
OT-CAR (20)	14	18	14	17	14	19	14	19
Σ	51	83	50	80	47	52	49	54

Domain	CPU-Time (seconds)							
	POLY				EXP			
	POLY ₀	POLY ₁	POLY ₂	POLY ₃	EXP ₀	EXP ₁	EXP ₂	EXP ₃
ROVER (20)	64.7	64.6	69.1	60.5	15.0	15.6	13.7	12.8
LIN-CAR (10)	10.7	7.0	7.2	8.3	8.8	7.5	5.0	8.1
LIN-GEN (10)	114.8	28.7	107.7	35.0	87.3	76.3	97.1	82.1
UTC (10)	—	90.6	—	73.0	—	—	—	—
BAXTER (20)	894.1	17.9	637.0	2.9	—	—	—	—
OT-CAR (20)	61.4	15.8	56.8	12.6	7.3	5.8	5.6	4.5

Domain	Makespan							
	POLY				EXP			
	POLY ₀	POLY ₁	POLY ₂	POLY ₃	EXP ₀	EXP ₁	EXP ₂	EXP ₃
ROVER (20)	522.2	522.2	522.2	522.2	522.2	522.2	522.2	522.2
LIN-CAR (10)	16.1	16.1	16.1	16.1	14.9	14.9	14.9	14.9
LIN-GEN (10)	1006.0	1006.0	1006.0	1006.0	1010.0	1010.0	1010.0	1010.0
UTC (10)	—	54.6	—	54.6	—	—	—	—
BAXTER (20)	5.0	5.0	5.0	5.0	—	—	—	—
OT-CAR (20)	42.1	42.1	41.2	41.2	24.5	24.5	23.8	23.8

Domain	Evaluated Nodes (1000)							
	POLY				EXP			
	POLY ₀	POLY ₁	POLY ₂	POLY ₃	EXP ₀	EXP ₁	EXP ₂	EXP ₃
ROVER (20)	14.7	14.7	14.7	14.7	8.7	8.7	8.7	8.7
LIN-CAR (10)	2.6	2.6	2.6	2.6	0.4	0.4	0.4	0.4
LIN-GEN (10)	11.1	11.1	11.1	11.1	2.1	2.1	2.0	2.0
UTC (10)	—	131.1	—	131.1	—	—	—	—
BAXTER (20)	0.2	0.1	0.2	0.1	—	—	—	—
OT-CAR (20)	122.1	122.1	92.0	92.0	4.9	4.9	5.6	5.6

Table 2: Domain by domain performance achieved by METRIC-FF when run with the presented translations, with $\delta = 1$, and optimisations. Results are presented in terms of coverage (number of solved instances), average CPU-Time, average quality (makespan), and average number of nodes expanded during the search process. Averages are calculated considering instances solved by all approaches. “—” indicates that no instances can be considered for the average calculation.

noticing that these domains are not necessarily leading to instances that are easier to be solved than those domains that include a larger number of processes and events – in fact, there is no direct relationship between the size and complexity of a problem to be solved.

On compact instances, the results presented in Table 2 indicate that the use of the EXP translation can generally lead to the best performance in terms of the number of evaluated nodes, CPU-Time, and quality of the generated solutions. Overall, the quality of computed plans tends to be similar, regardless of the translation used. The only notable exception

is OVERTAKING-CAR, where the EXP translations better support the planning process of METRIC-FF.

With regards to the optimisations, results in Table 2 suggest that POLY₁ is the best optimisation for POLY and that EXP₃ is the best combination for the EXP translation. Those optimisations are adopted for the remainder of the experimental analysis.

5.4 Results Contextualisation

We are now in the best position to contextualise the results achieved by the proposed translations using the selected planning engine, METRIC-FF, with the direct PDDL+ representation using different PDDL+ planning engines. Table 3 shows the achieved results in terms of the number of solved problems, by the considered planning approaches on the benchmark domains. It is also worth remarking that some of the PDDL+ planning engines required the models to be modified in order to generate a solution: those are indicated using “+”. The presented results highlight that the proposed translations are effective in supporting the use of PDDL2.1 planning engines for solving complex hybrid planning problems. In fact, the use of the proposed translations allows a PDDL2.1 planning engine to even outperform native PDDL+ engines. This suggests that the proposed translations can foster the exploitation of PDDL+ in real-world applications, by extending the pool of domain-independent planning engines.

Figure 2 gives some insights into the CPU-Time needed by the considered systems to solve the benchmark problems. All the approaches are able to quickly solve a large number of considered instances. When using the POLY translation, the PDDL2.1 planning engine is able to solve approximately 80 instances in less than 100 CPU-Time seconds. Notably, the curve of METRIC-FF with POLY does not flatten as quickly as others, suggesting that the combination of models and planning engine can effectively tackle also challenging instances, that require a significant amount of CPU-Time to be devoted to the search space exploration.

With regards to the quality of the generated plans, measured as makespan, we did not observe any significant overall difference between plans generated using the PDDL+ or the PDDL2.1 models. For a more extensive comparison, we used ENHSP with A (h^{blind}) heuristics (to ensure systematic exploration), which is equivalent to UPMURPHI in terms of quality of generated solutions. In particular, the comparison between the makespans of the plans found with A (h^{blind}) and METRIC-FF, used with different encodings and all optimisations, shows that there are some domains, such as SOLAR-ROVER and UTC, where plans with the same makespan are found. In a few other domains, such as BAXTER, LINEAR-CAR and OVERTAKING-CAR, METRIC-FF finds plans with longer makespans w.r.t. A (h^{blind}). For instance, in LINEAR-CAR, A (h^{blind}) finds solutions of average quality 11, compared to 14.9 of the best combination of METRIC-FF. In LINEAR-GENERATOR, both POLY and EXP find plans slightly longer (less than 0.1%) than A (h^{blind}). In general, it emerges that POLY, although it allows obtaining the best coverage, produces in some cases plans with longer makespan when compared with those obtainable with EXP. However, when comparing the plans generated using translated and original PDDL+ models, quality seems to be more affected by the planning approach exploited by the engine, rather than by the use of a specific formulation.

Domain	METRIC-FF		DiNo	ENHSP20	SMTPLAN	UPMURPHI
	POLY ₁	EXP ₃				
ROVER (20)	19	20	20 ⁺	5	19	4
LIN-CAR (10)	10	10	10 ⁺	10	10 ⁺	10
LIN-GEN (10)	10	5	10 ⁺	10	10 ⁺	1
UTC (10)	7	0	0	7	0	1
BAXTER (20)	19	0	7	17	8	12
OT-CAR (20)	18	19	0	19	0	5
Σ	83	54	47	68	47	33

Table 3: Number of problems solved by the considered planning approaches. Between brackets, we indicate the number of problem instances considered for each domain. POLY and EXP are used to indicate that, respectively, the polynomial or the exponential translation has been used, in both cases with $\delta = 1$. “+” denotes that the reported result refers to a variant of the domain model we considered, modified to allow the specific engine to reason upon it. Bold indicates the best results.




Figure 2: Total number of instances solved by each of the considered planning approaches, over time.

6. Related Work

A range of techniques have been introduced to support the reasoning on PDDL+ instances employing a translation. Balduccini, Magazzeni, Maratea, and Leblanc (2017) proposed an approach for translating a given PDDL+ instance in a Constraint ASP instance (Baselice, Bonatti, & Gelfond, 2005), but the process has to be done manually by an expert of the field. A number of approaches have been introduced to translate PDDL+ instances into Satisfiability Modulo Theories (SMT) (Barrett & Tinelli, 2018) problems (Bryce, Gao,

Musliner, & Goldman, 2015; Scala, Ramírez, Haslum, & Thiébaux, 2016; Cashmore et al., 2020) or a mix of linear programming and SAT instances (Shin & Davis, 2005). These translations differ from each other in the way the compilation is carried on; some of them make use of an intermediate translation step into hybrid automata (Bryce et al., 2015; Heinz, Wehrle, Bogomolov, Magazzeni, Greitschus, & Podelski, 2019; Cashmore et al., 2020), other ones perform a more direct translation (Shin & Davis, 2005). We follow on these lines, yet, as we map the problem into another planning problem, we are not required to provide an upper limit on the plan length as these approaches do. It is indeed the planning engine that needs to figure out the length of the plan complying with the problem constraints. Conversely, all the above translations require the ability to anticipate a maximum number of time points, which is a difficulty shared with simpler forms of planning too, such as SAT-based planning (Kautz & Selman, 1992; Rintanen, Heljanko, & Niemelä, 2006). Moreover, as we build PDDL2.1 formulations, we can in principle exploit all heuristics and search guidance mechanisms developed for PDDL2.1 problems (Hoffmann, 2003; Scala et al., 2016, 2020). We indeed do so indirectly by experimenting with all those planners that do implement such heuristics. Yet, as the number of planning engines that supports conditional effects effectively is very limited, we could only exploit METRIC-FF to a large extent. We highlight that this is not a limit of our translations, and believe that with the addition of newer and more expressive planning engines our translations can benefit the solving of PDDL+ even more.

Related to some extent is also the work by Coles and Coles (2014), which introduced a translation between a non-discretised PDDL+ problem instance and a PDDL2.1 temporal continuous instance, but showed that such way does not lead to models that are suitable for PDDL2.1 planning engines. Our approach is instead aiming at level 2 of the PDDL2.1 language. The language that we are targeting does not have a notion of time, with the result that constructing a planning engine supporting it is much easier than one that needs to provide native support for temporal reasoning.

A different line of work in automated planning focuses on reformulating models, to make them more amenable for planning engines, but without involving a translation to a different language or formalism. With regards to PDDL+ models, Franco, Vallati, Lindsay, and McCluskey (2019) introduced a technique for minimising the ground size of PDDL+ planning instances by reducing the arity of sparse predicates, i.e., predicates with a very large number of possible groundings, out of which very few are actually exploited in the planning problems.

There is an interesting parallel with compilations devised for classical planning models (Nebel, 2000; Gazen & Knoblock, 1997). In particular, our exponential translation anticipates the possible contexts a system is in much as the exponential encoding by (Gazen & Knoblock, 1997) compiles away conditional effects, while our polynomial translation captures the semantics of processes unrolling them into several actions, much as (Nebel, 2000) proposes to simulate the execution of conditional effects. As these two approaches have contributed to the discovery of several techniques and heuristics for classical planning (Haslum, 2013; Röger, Pommerening, & Helmert, 2014), we believe that our schemata can do the same for the much more involved case of PDDL+.

Finally, it is worth noting that our work provides a slightly alternative, more direct formalisation of the PDDL+ semantics. In fact, in the article presenting PDDL+ (Fox & Long,

2006), the authors resorted to a well-established formalism, known as Hybrid Automata (HA) (Henzinger, 1996), to explain the semantics of the proposed planning language. HAs are typically used for modelling systems having both a logical and a physical part where the latter is characterised by continuous dynamics. Bogomolov, Magazzeni, Podelski, and Wehrle (2014) further refine this translation into HAs by exploiting the more standard semantics of HA with specific attention devoted to must transitions prescribed by events and processes (Bogomolov et al., 2015). This line of research enables the use of tools specifically designed by the model-checking community, such as the SPACEEX model checker (Frehse et al., 2011).

Within the PDDL+ literature there has not been a systematic study of what it means to find an optimal plan whereas PDDL2.1 allows the definition of customised metrics to characterise the quality of the plans. Chen, Williams, and Fan (2021) address this problem by providing an encoding for translating a linear hybrid system into a Mixed Integer Linear Program (MILP), which can be tackled in an optimal way by using a MILP optimiser. Our work partially addresses this issue in the context of PDDL+ under discrete semantics, as the translation into numeric tasks, combined with a simple cost function that is incremented by δ when the time advances (in POLY when the *start* action is performed and in EXP when the *SIM* action is performed), would allow one to find optimal plans in terms of makespan if they were used in a search scheme that guarantees to demonstrate the optimality of a solution. Further studies are needed in this direction.

Our work grounds on the seminal paper by Shin and Davis (2005), in particular as far as it is concerned by the adopted temporal ontology that is based on the notion of *superdense time* (Maler, Manna, & Pnueli, 1991). A superdense time model extends the real-valued timeline with additional information necessary to model the ordering of multiple simultaneous transitions. The extended timeline model provided by Shin and Davis (2005) allows us to capture when, informally speaking, something happens, i.e., an action, event or process starts or ends but time does not flow. Previous usages of this kind of temporal model can be traced back to TLPLAN (Bacchus & Ady, 2001) and OPTOP (McDermott, 2003).

Batusov and Soutchanski (2019) recently propose an alternative logical semantic of PDDL+ which extends *Situation Calculus* (SC), i.e., a logical formalism to represent dynamic domains, giving it a continuous component inspired by the HA interpretation. The target formalism of this mapping is an extension of a particular case of SC, namely *Basic Action Theory* (Reiter, 2001), which is a convenient representation for addressing reasoning problems. There is finally an interesting analogy on how hybrid automata are formalised in terms of timed transition systems (Henzinger, 1996). Indeed, our translations can be understood as a way to direct reasoning over the timed transition system and therefore use a *simpler* planner. In future work, we aim at studying this interpretation in deeper detail with the hope that this would further a prolific cross-fertilisation between techniques developed for planning models and hybrid automata.

7. Discussion

In this section, we frame the scope of this work and we discuss potential challenges that discretisation poses.

One of the main aims of this work is to show that solving a PDDL+ problem under a discrete interpretation of the timeline is equivalent to solving a PDDL2.1 problem without durative actions. This necessitated the formalisation of discrete PDDL+ according to a discretisation step $\delta \geq \mathbb{Q}_{>0}$, which we assume to be an external parameter of the problem. This contribution, to the best of our knowledge, is the first attempt to formalise a discrete semantics for PDDL+ and we believe it is useful to the extent that it provides us with a formal framework for evaluating the correctness of those PDDL+ planning engines that natively work on discretised PDDL+.

In a wider sense, one may argue about the practical value of discretised models. A major practical benefit lies in the fact that discretised models allow planning engines to solve challenging and complicated tasks, otherwise impossible to solve by the current state of the art. There is of course the open question about the relation between PDDL+ and its discretised version, which however is beyond the scope of this work. We treat the two problems as two inherently different problems. It shall be noted that not all the possible discretisations have equal value concerning the problem to be solved. There is a trade-off to consider between large and small δ values; a very large δ can drastically reduce the complexity of the planning process, at the cost of an approximation of numeric dynamics that can lead the model to diverge too much from the ideal one, thus making it of little interest in practical terms. An extremely small value of δ increases the complexity of the instance to be solved, not guaranteeing at the same time the validity of solutions in the continuous settings, but allowing a more accurate approximation of the dynamics of the problem. The search for a suitable δ can be done by validating the solutions found towards smaller δ as we did in a recent work (Percassi, Scala, & Vallati, 2022); in this way, it would be possible to compute solutions with large δ_s (which stands for *search* δ) by conducting a more lightweight search and validate them effectively with an arbitrarily small δ_v (which stands for *validation* δ) which is sufficiently reliable according to the practitioner.

It is also worth noting that the discretisation of a PDDL+ problem may cause a tightening of its solution space because all the plans having actions executed in $\mathbb{R} \cap \delta \cdot \mathbb{Z}$ are expunged from the pool of possible solutions. However, there are problems for which even admitting arbitrarily small δ , would remain unsolved under discrete semantics. For example, suppose to have LINEAR-CAR problem in which the car has to raise a flag at the following distance $d = \frac{1}{2}$ and $d = 1$. In a continuous sense this would mean executing an action, e.g., *raiseFlag*, at time $t = \frac{1}{2}$ and $t = 1$. It follows that, since the numeric variables in PDDL2.1 problems take values in \mathbb{Q} , and the actions are timestamped in \mathbb{Q} , it is not possible to find a rational δ , even arbitrarily small, to solve this problem continuously. It is worth mentioning that Fox and Long (2003) discuss this possibility, suggesting that it is not possible to obtain an arbitrary precision of the plans, or over other numeric quantity, but it is possible to pragmatically validate them by admitting a numeric tolerance. Such a problem is known also in the context of HA (Henzinger & Raskin, 2000).

8. Conclusion

Hybrid PDDL+ models are amongst the most advanced models of systems and the resulting problems are notoriously difficult for planning engines to cope with. To deepen the understanding of PDDL+, and to support the solvability of PDDL+ instances, in this paper

we introduced two translations from time-discretised PDDL+ to PDDL2.1 (level 2). The exponential translation leads to a numeric planning problem which is exponentially larger than the initial PDDL+ but preserves the number of discrete transitions. The polynomial translation instead leads to a smaller formulation but requires more transitions to generate a solution. We also presented two optimisations for the proposed translation schemata, that aim at reducing the size of the translated models while preserving the soundness and completeness of the approach. The optimisations exploit information about the structure of the considered planning instance to avoid unnecessary checks on actions and events.

Our experimental analysis demonstrated the usefulness of the introduced translations and optimisations in unlocking the exploitation of PDDL2.1 planning engines to solve challenging PDDL+ instances. The introduced optimisations make instances more amenable for domain-independent PDDL2.1 planning engines. In particular, the use of the optimised polynomial translation allows a PDDL2.1 planning engine to outperform a state-of-the-art PDDL+ planning engine across a wide range of benchmark domains. Summarising, the proposed translations can unlock the use of PDDL2.1 planning engines for tackling hybrid PDDL+ problems, with the clear advantage of significantly expanding the number of approaches that can be used to solve a problem instance.

In future, we plan to explore incomplete translations, where a trade-off can potentially be found between completeness and the size of the resulting PDDL2.1 instances. We are interested in incorporating the introduced translations into existing planning engines, possibly targeting the grounding step to minimise overhead (Scala & Vallati, 2021). Finally, we are interested in investigating potential synergies between the proposed translations and well-known reformulation approaches such as macro-actions, to generate PDDL2.1 models that are more suitable to domain-independent planning engines.

Acknowledgements

Francesco Percassi and Mauro Vallati are supported by the UKRI Future Leaders Fellowship [grant number MR/T041196/1]. Enrico Scala has been partially supported by AIPlan4EU, a project funded by EU Horizon 2020 research and innovation programme under GA n. 101016442.

References

- Bacchus, F., & Ady, M. (2001). Planning with Resources and Concurrency: A Forward Chaining Approach In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001*, pp. 417–424.
- Balduccini, M., Magazzeni, D., Maratea, M., & Leblanc, E. C. (2017). CASP Solutions for Planning in Hybrid Domains *Theory and Practice of Logic Programming*, 17(4), 591–633.
- Barrett, C., & Tinelli, C. (2018). Satisfiability Modulo Theories In *Handbook of model checking*, pp. 305–343. Springer.
- Baselice, S., Bonatti, P. A., & Gelfond, M. (2005). Towards an Integration of Answer Set and Constraint Solving In *Proceedings of the Twenty-First International Conference*

- on *Logic Programming, ICLP 2005*, pp. 52–66.
- Batusov, V., & Soutchanski, M. (2019). A Logical Semantics for PDDL+ In *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2019*, pp. 40–48.
- Benton, J., Coles, A. J., & Coles, A. (2012). Temporal Planning with Preferences and Time-Dependent Continuous Costs In *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling, ICAPS 2012*, pp. 2–10.
- Bertolucci, R., Capitanelli, A., Maratea, M., Mastrogiovanni, F., & Vallati, M. (2019). Automated Planning Encodings for the Manipulation of Articulated Objects in 3D with Gravity In *Proceedings of the Eighteenth International Conference of the Italian Association for Artificial Intelligence, AI*IA 2019*, pp. 135–150.
- Bogomolov, S., Magazzeni, D., Minopoli, S., & Wehrle, M. (2015). PDDL+ Planning with Hybrid Automata: Foundations of Translating Must Behavior In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015*, pp. 42–46.
- Bogomolov, S., Magazzeni, D., Podelski, A., & Wehrle, M. (2014). Planning as Model Checking in Hybrid Domains In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2014*, pp. 2228–2234.
- Bonassi, L., Gerevini, A. E., Percassi, F., & Scala, E. (2021). On Planning with Qualitative State-Trajectory Constraints in PDDL3 by Compiling them Away In *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling, ICAPS 2021*, pp. 46–50.
- Bryce, D., Gao, S., Musliner, D. J., & Goldman, R. P. (2015). SMT-Based Nonlinear PDDL+ Planning In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI 2015*, pp. 3247–3253. AAAI Press.
- Cashmore, M., Magazzeni, D., & Zehtabi, P. (2020). Planning for Hybrid Systems via Satisfiability Modulo Theories *Journal of Artificial Intelligence Research*, 67, 235–283.
- Chen, J., Williams, B. C., & Fan, C. (2021). Optimal Mixed Discrete-Continuous Planning for Linear Hybrid Systems In *Proceedings of the Twenty-Fourth International Conference on Hybrid Systems: Computation and Control, HSCC 2021*, pp. 1–12.
- Coles, A. J., & Coles, A. I. (2014). PDDL+ Planning with Events and Linear Processes In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS 2014*, pp. 74–82.
- Cooper, M. C., Maris, F., & Régnier, P. (2010). Compilation of a High-level Temporal Planning Language into PDDL 2.1 In *Proceedings of the Twenty-Second IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2010*, pp. 181–188.
- Fox, M., Howey, R., & Long, D. (2005). Validating Plans in the Context of Processes and Exogenous Events In *Proceedings of the Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference*, pp. 1151–1156.

- Fox, M., & Long, D. (2003). PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains *Journal of Artificial Intelligence*, 20, 61–124.
- Fox, M., & Long, D. (2006). Modelling Mixed Discrete-Continuous Domains for Planning *Journal of Artificial Intelligence Research*, 27, 235–297.
- Franco, S., Vallati, M., Lindsay, A., & McCluskey, T. L. (2019). Improving Planning Performance in PDDL+ Domains via Automated Predicate Reformulation In *Proceedings of the Nineteenth Conference of Computational Science, ICCS 2019*, pp. 491–498.
- Frehse, G., Guernic, C. L., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., & Maler, O. (2011). SpaceEx: Scalable Verification of Hybrid Systems In *Proceedings of the Twenty-Third International Conference on Computer Aided Verification, CAV 2011*, pp. 379–395.
- Gazen, B. C., & Knoblock, C. A. (1997). Combining the Expressivity of UCPOP with the Efficiency of Graphplan In *Proceedings of Recent Advances in AI Planning: Fourth European Conference on Planning, ECP 1997*, pp. 221–233.
- Gerevini, A., Saetti, A., & Serina, I. (2008). An approach to efficient planning with numerical fluents and multi-criteria plan quality *Artificial Intelligence*, 172(8-9), 899–944.
- Gigante, N., Micheli, A., Montanari, A., & Scala, E. (2022). Decidability and complexity of action-based temporal planning over dense time *Artificial Intelligence*, 307, 103686.
- Grastien, A., & Scala, E. (2020). CPCES: A planning framework to solve conformant planning problems through a counterexample guided refinement *Artificial Intelligence*, 284, 103271.
- Haslum, P. (2013). Optimal Delete-Relaxed (and Semi-Relaxed) Planning with Conditional Effects In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI 2013*, pp. 2291–2297.
- Heinz, A., Wehrle, M., Bogomolov, S., Magazzeni, D., Greitschus, M., & Podelski, A. (2019). Temporal Planning as Refinement-Based Model Checking In *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2019*, pp. 195–199.
- Helmert, M. (2009). Concise finite-domain representations for PDDL planning tasks *Artificial Intelligence*, 173(5-6), 503–535.
- Henzinger, T. A. (1996). The Theory of Hybrid Automata In *Proceedings of the Eleventh Annual IEEE Symposium on Logic in Computer Science, LICS 1996*, pp. 278–292.
- Henzinger, T. A., & Raskin, J. (2000). Robust Undecidability of Timed and Hybrid Systems In *Third International Workshop on Hybrid Systems: Computation and Control, HSCC 2000*, Vol. 1790, pp. 145–159.
- Hoffmann, J. (2003). The Metric-FF Planning System: Translating “Ignoring Delete Lists” to Numeric State Variables *Journal of Artificial Intelligence Research*, 20, 291–341.
- Kautz, H. A., & Selman, B. (1992). Planning as Satisfiability In *Proceedings of the Tenth European conference on Artificial intelligence, ECAI 1992*, pp. 359–363.
- Maler, O., Manna, Z., & Pnueli, A. (1991). From Timed to Hybrid Systems In *Proceedings of the Real-Time: Theory in Practice, REX Workshop*, Vol. 600, pp. 447–484.

- McCluskey, T. L., & Porteous, J. M. (1997). Engineering and compiling planning domain models to promote validity and efficiency *Artificial Intelligence*, 95(1), 1–65.
- McCluskey, T. L., & Vallati, M. (2017). Embedding Automated Planning within Urban Traffic Management Operations In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017*, pp. 391–399. AAAI Press.
- McCluskey, T. L., Vaquero, T. S., & Vallati, M. (2017). Engineering Knowledge for Automated Planning: Towards a Notion of Quality In *Proceedings of the Knowledge Capture Conference, K-CAP 2017*, pp. 14:1–14:8. ACM.
- McDermott, D. V. (2003). Reasoning about Autonomous Processes in an Estimated-Regression Planner In *Proceedings of the Thirteenth International Conference on Automated Planning and Scheduling, ICAPS 2003*, pp. 143–152. AAAI.
- Meurer, A., Smith, C. P., Paprocki, M., Čertík, O., Kirpichev, S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K., Singh, S., Rathnayake, T., Vig, S., Granger, B. E., Muller, R. P., Bonazzi, F., Gupta, H., Vats, S., Johansson, F., Pedregosa, F., Curry, M. J., Terrel, A. R., Roučka, v., Saboo, A., Fernando, I., Kulal, S., Cimrman, R., & Scopatz, A. (2017). SymPy: symbolic computing in Python *PeerJ Computer Science*, 3, e103.
- Nebel, B. (2000). On the Compilability and Expressive Power of Propositional Planning Formalisms *Journal of Artificial Intelligence Research*, 12, 271–315.
- Palacios, H., & Geffner, H. (2009). Compiling Uncertainty Away in Conformant Planning Problems with Bounded Width *Journal of Artificial Intelligence Research*, 35, 623–675.
- Penna, G. D., Magazzeni, D., & Mercorio, F. (2012). A universal planning system for hybrid domains *Applied Intelligence*, 36(4), 932–959.
- Percassi, F., & Gerevini, A. E. (2019). On Compiling Away PDDL3 Soft Trajectory Constraints without Using Automata In *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2019*, pp. 320–328. AAAI Press.
- Percassi, F., Scala, E., & Vallati, M. (2021). Translations from Discretised PDDL+ to Numeric Planning In *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling, ICAPS 2021*, pp. 252–261. AAAI Press.
- Percassi, F., Scala, E., & Vallati, M. (2022). The Power of Reformulation: From Validation to Planning in PDDL+ In *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling, ICAPS 2022*, pp. 288–296. AAAI Press.
- Piotrowski, W. M., Fox, M., Long, D., Magazzeni, D., & Mercorio, F. (2016). Heuristic planning for hybrid systems In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pp. 4254–4255. AAAI Press.
- Reiter, R. (2001). *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT press.
- Rintanen, J., Heljanko, K., & Niemelä, I. (2006). Planning as satisfiability: parallel plans and algorithms for plan search *Artificial Intelligence*, 170(12-13), 1031–1080.

- Röger, G., Pommerening, F., & Helmert, M. (2014). Optimal Planning in the Presence of Conditional Effects: Extending LM-Cut with Context Splitting In *Proceedings of the Twenty-First European Conference on Artificial Intelligence, ECAI 2016*, pp. 765–770.
- Scala, E., Haslum, P., Thiébaux, S., & Ramírez, M. (2016). Interval-Based Relaxation for General Numeric Planning In *Proceedings of the Twenty-Second European Conference on Artificial Intelligence, ECAI 2016*, Vol. 285, pp. 655–663.
- Scala, E., Haslum, P., Thiébaux, S., & Ramírez, M. (2020). Subgoaling Techniques for Satisficing and Optimal Numeric Planning *Journal Artificial Intelligence Research*, 68, 691–752.
- Scala, E., Ramírez, M., Haslum, P., & Thiébaux, S. (2016). Numeric Planning with Disjunctive Global Constraints via SMT In *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling, ICAPS 2016*, pp. 276–284. AAAI Press.
- Scala, E., & Vallati, M. (2021). Effective grounding for hybrid planning problems represented in PDDL+ *The Knowledge Engineering Review*, 36.
- Shin, J., & Davis, E. (2005). Processes and continuous change in a SAT-based planner *Artificial Intelligence*, 166(1-2), 194–253.
- Taig, R., & Brafman, R. I. (2013). Compiling Conformant Probabilistic Planning Problems into Classical Planning In *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling, ICAPS 2013*. AAAI Press.
- Vallati, M., Chrapa, L., McCluskey, T. L., & Hutter, F. (2021). On the Importance of Domain Model Configuration for Automated Planning Engines *Journal of Automated Reasoning*, 65(6), 727–773.
- Vallati, M., Magazzeni, D., Schutter, B. D., Chrapa, L., & McCluskey, T. L. (2016). Efficient Macroscopic Urban Traffic Models for Reducing Congestion: A PDDL+ Planning Approach In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI 2016*, pp. 3188–3194. AAAI Press.