

# Efficient Anytime Computation and Execution of Decoupled Robustness Envelopes for Temporal Plans

Michael Cashmore  

Strathclyde University, UK

Alessandro Cimatti  

Fondazione Bruno Kessler, Italy

Daniele Magazzeni  

Kings College London, UK

Andrea Micheli  

Fondazione Bruno Kessler, Italy

Parisa Zehtabi  

Kings College London, UK

---

## Abstract

One of the major limitations for the employment of model-based planning and scheduling in practical applications is the need of costly re-planning when an incongruence between the observed reality and the formal model is encountered during execution. Robustness Envelopes characterize the set of possible contingencies that a plan is able to address without re-planning, but their exact computation is expensive; furthermore, general robustness envelopes are not amenable for efficient execution.

In this paper, we present a novel, anytime algorithm to approximate Robustness Envelopes, making them scalable and executable. This is proven by an experimental analysis showing the efficiency of the algorithm, and by a concrete case study where the execution of robustness envelopes significantly reduces the number of re-plannings.

**2012 ACM Subject Classification** Computing methodologies → Robotic planning

**Keywords and phrases** Temporal Planning, Robustness Envelopes

**Acknowledgements** The open access publication of this article was supported by the Alpen-Adria-Universität Klagenfurt, Austria.

## 1 Introduction

When planning and scheduling techniques are employed in practical applications, one of the major problems is the need for on-line re-planning when the observed contingencies are not aligned with the ones that were considered at planning time. These situations are common, because it is arguably impossible to predict the entire range of situations an autonomous system can encounter, especially when the planning domain encompasses time and temporal constraints. Unfortunately, re-planning can be costly in terms of time, and computational resources can be scarce on-board, so limiting the use of re-planning is very important for practical purposes. In principle, it is also possible to continue with the execution of a plan even when the observed contingencies are unexpected, optimistically hoping for a successful completion. However, this approach offers no formal guarantee, and is prone to the risk of continuing execution of a plan that is bound to fail.

Several approaches have been proposed in the literature to address this problem (see [13] for a survey focused on robotics). Some authors propose to post-process plans and generalize them relying on the scheduling constraints that are relevant for execution [17, 15, 10]. Another line of research focuses on the creation of “least commitment plans”, i.e. plans that are left

44 partially open by the planner so that the execution can be adapted to some variation in the  
 45 contingencies [11, 20, 9, 4, 19]. Others tackled the idea of transforming temporal plans with  
 46 no adaptability into flexible plans [7]. Finally, one can explicitly model the uncertainties in  
 47 the planning problem and construct a plan that offers formal guarantees with respect to  
 48 such a model. Examples include Conformant and Contingent Planning [12], Probabilistic  
 49 Planning [14] and Strong Temporal Planning with Uncertain Durations [5] that considers  
 50 temporal uncertainty in the durations of actions.

51 Recently, *Robustness Envelopes* (REs) have been proposed to overcome several limitations  
 52 of the approaches mentioned above. REs formally capture the possible contingencies that  
 53 a given temporal plan, obtained by planning in a *deterministic domain*, can deal with,  
 54 without having to re-plan [2]. REs are regions defined over a set of numeric parameters  
 55 that represent possible contingencies, and contain all the parameter valuations ensuring  
 56 plan validity. In general, REs may be non-convex, and can express dependencies between  
 57 the parameters. However, the technique proposed in [2] has two main drawbacks limiting  
 58 its practical applicability. First, the exact computation of REs is extremely expensive: the  
 59 proposed approach is doubly exponential in the size of the planning problem. Second, REs in  
 60 their general form are not suited for efficient execution: the dependencies among parameters  
 61 might require run-time reasoning.

62 In this paper, we overcome these limitations, achieving scalability and executability.  
 63 We focus on *Decoupled Robustness Envelopes* (DREs), i.e. hyper-rectangular REs where  
 64 the dependencies among parameters are not present, and are thus much easier to execute.  
 65 Our first contribution is a novel and scalable algorithm for computing DREs as sound  
 66 approximations of REs. A sound approximation ensures that every point within the DRE  
 67 belongs within the RE. The algorithm is anytime, and proceeds by incrementally under-  
 68 approximating the RE with increasingly large DREs. The algorithm can be stopped at  
 69 any time, providing a meaningful result already amenable to start execution. In its general  
 70 formulation, the RE for a given plan is naturally modeled as a quantified first order formula  
 71 in the theory of Linear Real Arithmetic. Our algorithm does not need to precisely compute  
 72 the quantifier-free description of the RE (which requires an expensive step of quantifier  
 73 elimination, and is ultimately responsible for the inefficiency demonstrated in [2]). Rather, it  
 74 starts from a degenerate DRE consisting of a single point, and progressively tries to enlarge  
 75 it along different dimensions, checking if each extension is contained in the RE, until a given  
 76 precision is reached. The algorithm relies on *quantifier-free* queries to a Satisfiability Modulo  
 77 Theory [1] solver.

78 Our second contribution is to demonstrate the practical use of DREs in a robotic executor,  
 79 extending the classical flow from planning to execution to re-planning, as follows. First,  
 80 a plan is generated from a deterministic model using temporal planning technologies, and  
 81 transformed into a Simple Temporal Network (STN) formulation [6]; at this point, we  
 82 parametrize the durations of some of the actions in the plan and/or the consumption rates  
 83 in the domain specification. DREs are then computed for the introduced parameters and  
 84 passed to the executor. In turn, the dispatching of the actions in the STN plan begins and  
 85 continues until one observed duration or consumption rate happens to be outside of the DRE.  
 86 At this point, the executor detects that the plan is no longer guaranteed to succeed, and  
 87 re-planning is triggered.

88 The proposed approach was implemented in the ROSPlan [3] framework, and experiment-  
 89 ally evaluated along two directions. The algorithm for DRE generation was compared against  
 90 the base line in [2], demonstrating orders-of-magnitude improvements compared to the exact  
 91 computation of REs, and the ability to deal with a much larger number of parameters. The

92 overall execution loop has been evaluated on a family of concrete case studies in a logistic  
 93 domain, showing that the use of DREs, compared to the optimistic executor in ROSPlan,  
 94 significantly reduces the number of re-plannings and improves the execution success-rate.

## 95 **2** Background

96 We consider planning problems expressed in the PDDL 2.1 [8] temporal planning language;  
 97 for the sake of brevity we do not report the full syntax and semantics of such planning  
 98 problems, but we directly introduce the parametrized planning problem idea adapted from [2].

99 ► **Definition 1.** A *parametrized planning problem*  $\mathcal{P}_\Gamma$  is a tuple  $\langle \Gamma, \mathcal{P} \rangle$ , where  $\Gamma$  is a  
 100 finite set of real-valued parameters  $\{\gamma_1, \dots, \gamma_n\}$  and  $\mathcal{P}$  is a PDDL 2.1 planning problem in  
 101 which action conditions, action effects, goals and initial states can all contain parameters.

102 Intuitively, symbols (from a known set  $\Gamma$ ) can be used in expressions where real-typed  
 103 constants are usually allowed.

104 As customary in many cases of plan execution, we use plans expressed as Simple Temporal  
 105 Networks (STN) [6]. An STN plan is a set of constraints of the form  $t_i - t_j \leq k$  where  $t_i$  and  
 106  $t_j$  are time points linked to action happenings (i.e. either the start or the end of an action  
 107 instance in the plan) and  $k \in \mathbb{Q}$ . In addition, we allow parameters in the plan specification  
 108 by generalizing the notion of STN plans.

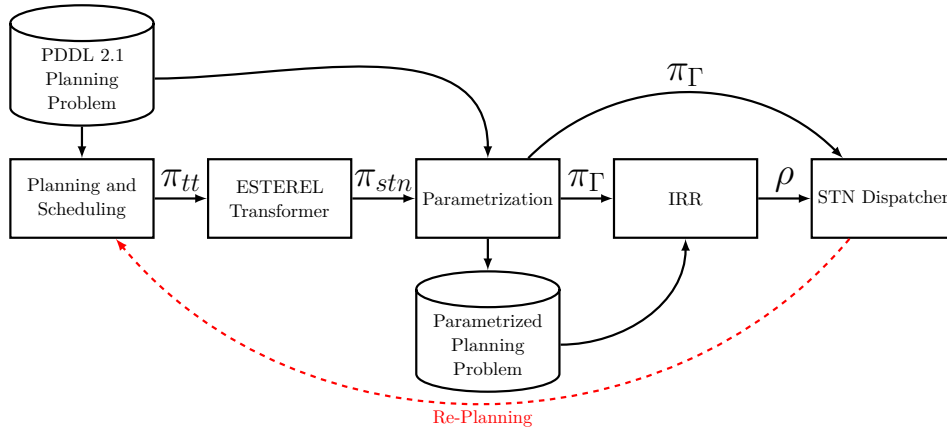
109 ► **Definition 2.** A *parametrized STN plan*  $\pi_\Gamma$  for a parametrized planning problem  
 110  $\mathcal{P}_\Gamma \doteq \langle \Gamma, \mathcal{P} \rangle$  is an STN Plan where some constraints are in the form  $t_i - t_j \leq \gamma_i$  where  $t_i$   
 111 and  $t_j$  are time-points of the STN plan and  $\gamma_i \in \Gamma$ .

We define the Robustness Envelope (RE) for a parametrized problem and plan as the set of possible values for the parameters that make the plan valid when the symbols are substituted with values in the plan and problem specifications. In order to compute the RE, [2] defines a set of logical formulae that characterize the RE and use quantifier elimination techniques (e.g. [18]) to explicitly construct the region. The encoding is divided in three expressions: indicated as  $enc_{tn}^{\pi_\Gamma}$ ,  $enc_{eff}^{\pi_\Gamma}$  and  $enc_{proofs}^{\pi_\Gamma}$ . The formula  $enc_{tn}^{\pi_\Gamma}$  encodes the temporal constraints imposed by  $\pi_\Gamma$  limiting the possible orderings of time points. The formula  $enc_{eff}^{\pi_\Gamma}$  encodes the effects of each time point on the state variables, while  $enc_{proofs}^{\pi_\Gamma}$  encodes the validity properties of the plan, namely that the conditions of each executed action are satisfied, that the goal is reached, and that the  $\epsilon$ -separation constraint [8] imposed by PDDL 2.1 is respected. Then, let  $\bar{X}$  be the set of all the variables appearing in the formulae above except the parameter values, the RE is characterized by all the models of the following formula.

$$\exists \bar{X}. (enc_{tn}^{\pi_\Gamma} \wedge enc_{eff}^{\pi_\Gamma}) \wedge \forall \bar{X}. ((enc_{tn}^{\pi_\Gamma} \wedge enc_{eff}^{\pi_\Gamma}) \rightarrow enc_{proofs}^{\pi_\Gamma})$$

112 As observed in [2], any under-approximation of the RE gives sound information on  
 113 the contingencies in which the plan is guaranteed to be valid; in particular, a convenient  
 114 restriction for the representation and handling of REs is to associate a closed interval of  
 115 possible values to each parameter, defining an hyper-rectangle. If such hyper-rectangle is  
 116 contained in the RE, we have a "Decoupled Robustness Envelope" (DRE) that retains the  
 117 guarantees of the RE but avoids the complexity of inter-dependencies among parameters.

118 ► **Definition 3.** A *Decoupled Robustness Envelope* for a parametrized planning problem  
 119  $\mathcal{P}_\Gamma$  and STN plan  $\pi_\Gamma$  is a bound assignment  $\rho : \Gamma \rightarrow \mathbb{Q}_{>=0} \times \mathbb{Q}_{>=0}$ , such that any parameter  
 120 assignment  $v : \Gamma \rightarrow \mathbb{Q}_{>=0}$ , with  $l \leq v(\gamma) \leq u$  and  $\langle l, u \rangle \doteq \rho(\gamma)$ , is contained in the RE for  
 121  $\mathcal{P}_\Gamma$  and  $\pi_\Gamma$ .



■ **Figure 1** Overview of the proposed flow.


122 Note that many DREs are possible for a given problem and plan: it suffices that all the  
 123 assignments allowed by the DRE are points in the RE. In this paper, we elaborate on this  
 124 idea and propose an algorithm that incrementally builds DREs that are contained within  
 125 the unconstrained RE without paying the cost of explicitly computing the RE itself.

126 Finally, we highlight that given any two DREs  $\rho_1$  and  $\rho_2$  three cases are possible: either  
 127  $\rho_1$  is subsumed by  $\rho_2$  (i.e. for each parameter  $\gamma$ ,  $\rho_1(\gamma) \subseteq \rho_2(\gamma)$ ), or  $\rho_1$  subsumes  $\rho_2$ , or the  
 128 two DREs are incomparable. Hence, there is no absolute best DRE in general: we aim for a  
 129 DRE that is not subsumed by any other, but there can be multiple DREs with this property.

### 130 3 Execution Flow

131 The general idea we pursue in this paper is to exploit the information and the generalization  
 132 provided by the synthesis of REs to limit the number of re-plannings and increase the  
 133 success-rate in execution. In particular, we propose the flow from planning to execution  
 134 depicted in Figure 1. Starting from a planning problem formulation expressed in PDDL  
 135 2.1, we use any off-the-shelf temporal planner<sup>1</sup> to compute a timed sequence of actions that  
 136 reaches the goal from the initial state. We call this plan “time-triggered” (indicated with  
 137  $\pi_{tt}$ ) in the picture. This plan is not natively amenable for execution because it defines one  
 138 specific trace that does not allow any adaptability: it is extremely unlikely for a real system  
 139 to be perfectly controlled to satisfy a specific trace. Hence,  $\pi_{tt}$  needs to be converted in  
 140 a flexible, executable STN ( $\pi_{stn}$ ) by the ESTEREL transformer of ROSPlan. The usual  
 141 flow would pass this STN directly to the dispatcher for translating the plan actions into  
 142 commands for the robotic platform at the proper time. Instead, here we pre-process this  
 143 plan using REs in the hope of generalizing its applicability and reducing the number of  
 144 re-plannings. In particular, the STN plan is passed to a parametrization component that  
 145 re-reads the planning problem formulation and enriches it with parameters, generating a  
 146 Parametric Planning Problem and a parametric STN plan ( $\pi_{\Gamma}$ ). Those are the inputs for the

<sup>1</sup> Several existing PDDL planners are unable to generate flexible STNs either because of an implementation limitation or because the technique does not allow it (e.g. SAT-based planners). Our approach is able to generate DREs from these planners as well, and work in concert with existing algorithms for the execution of STNs.



■ **Figure 2** A graphical representation of IRR: starting from the parameter values from the original plan (depicted as the black point), IRR tries to construct increasingly better under-approximations (the colored rectangles) of the RE (the gray area), without actually computing it. Upon termination, each edge of the resulting DRE is guaranteed to be at most  $\beta$  apart from the border of the actual region.

147 computation of the RE. In our flow, for performance reasons and to avoid complex run-time  
 148 reasoning, instead of computing the exact, unconstrained RE, we use a novel algorithm,  
 149 called Incremental Rectangular-Robustification (IRR for short), that computes a DRE. The  
 150 algorithm is anytime, so that it is possible to retrieve unfinished computations and exploit  
 151 them in execution: in fact, any under-approximation of the final result retains the needed  
 152 properties of the RE. At this point, we pass the DRE ( $\rho$ ) together with the parametrized  
 153 STN plan to the STN dispatcher. We modified the dispatching algorithm to exploit the  
 154 information in the DRE to limit the re-plannings to situations where they are needed. In  
 155 particular, the dispatcher translates the actions, while checking that the observed values  
 156 for the parameters (being either action durations, resources or rates) fall within the bounds  
 157 imposed by  $\rho$ . If this is not the case, re-planning is needed and the whole flow is re-executed.

158 **Parametrization.** The first non-standard step highlighted in Figure 1 is the parametriza-  
 159 tion. In fact, there are multiple ways in which parameters can be added to a deterministic  
 160 temporal planning problem to characterize useful quantities for execution. In general, one  
 161 can parametrize any numeric quantity in the planning problem whose value might differ  
 162 from the environment in which the plan will be executed. In order to be useful for the STN  
 163 dispatcher, however, such quantities must be eventually observable (directly or indirectly).  
 164 Otherwise, it is impossible for the executor to check whether the RE is still satisfied or if  
 165 a re-planning is needed. In this paper experimentation, we focused on two such quantities,  
 166 namely the durations of actions and resource consumption rates. The former is a classical  
 167 source of uncertainty when temporal planning is employed in a robotics scenario, the latter  
 168 is another source of uncertainty that can perturbate the execution of a plan, for example  
 169 when the resource harvesting is not fully controllable (e.g. a solar panel yield depends on the  
 170 weather) or when the consumption is not fully predictable (e.g. the battery consumption is  
 171 very hard to precisely estimate as it depends on temperature, exact capacity and so on).

## 172 4 Incremental Rectangular-Robustification

173 We now present our novel algorithm for incrementally computing decoupled robustness  
 174 envelopes. We call this algorithm "Incremental Rectangular-Robustification" (IRR).

175 The idea behind the algorithm is to construct incrementally better hyper-rectangular  
 176 under-approximations of the RE for a given problem and plan. In fact, this constitutes a

---

**Algorithm 1** Incremental Rectangular-Robustification
 

---

```

1:  $enc_{valid} \leftarrow \text{QUANTIFIERELIMINATION}(\exists \bar{X}. enc_{tn}^{\pi_{\Gamma}} \wedge enc_{eff}^{\pi_{\Gamma}})$ 
2: function IRR( $\beta : \mathbb{Q}_{>0}$ )
3:    $R \leftarrow \{\gamma \rightarrow [\pi(\gamma), \pi(\gamma)] \mid \gamma \in \Gamma\}$ 
4:    $\Delta \leftarrow \{\gamma \rightarrow \max(\pi(\gamma) \times \omega_{\gamma}, \beta) \mid \gamma \in \Gamma\}$ 
5:    $\Theta \leftarrow \{\gamma \rightarrow \{\text{UB}, \text{LB}\} \mid \gamma \in \Gamma\}$ 
6:   while  $\exists \gamma \in \Gamma. \Delta(\gamma) \geq \beta$  do
7:      $\tilde{\gamma} \leftarrow \text{PICK}(\{\gamma \mid \gamma \in \Gamma \wedge \Theta(\gamma) \neq \emptyset \wedge \Delta(\gamma) \geq \beta\})$ 
8:      $\theta \leftarrow \text{PICK}(\Theta(\tilde{\gamma}))$ 
9:      $[l, u] \leftarrow R(\tilde{\gamma})$ 
10:    if  $\theta = \text{UB}$  then  $u \leftarrow (u + \Delta(\tilde{\gamma}))$  else  $l \leftarrow (l - \Delta(\tilde{\gamma}))$ 
11:     $R' \leftarrow \{\gamma \rightarrow R(\gamma) \mid \gamma \in \Gamma \wedge \gamma \neq \tilde{\gamma}\} \cup \{\tilde{\gamma} \rightarrow [l, u]\}$ 
12:    if CHECKINENVELOPE( $R'$ ) then  $R \leftarrow R'$ 
13:    else
14:       $\Theta(\tilde{\gamma}) \leftarrow \Theta(\tilde{\gamma}) \setminus \theta$ 
15:      if  $\Theta(\tilde{\gamma}) = \emptyset$  then
16:         $\Delta(\tilde{\gamma}) \leftarrow \Delta(\tilde{\gamma})/2$ ;  $\Theta(\tilde{\gamma}) \leftarrow \{\text{LB}, \text{UB}\}$ 
17:    return  $R$ 
18: function CHECKINENVELOPE( $R$ )
19:    $enc_R \leftarrow \bigwedge_{\gamma \in \Gamma, [l, u] = R(\gamma)} l \leq \tilde{\gamma} \wedge \tilde{\gamma} \leq u$ 
20:   if ISSAT( $enc_R \wedge \neg enc_{valid}$ ) then return false
21:   else
22:     return ISVALID( $(enc_{tn}^{\pi_{\Gamma}} \wedge enc_{eff}^{\pi_{\Gamma}} \wedge enc_R) \rightarrow enc_{proofs}^{\pi_{\Gamma}}$ )
```

---

177 direct way of computing a DRE by generate-and-test. The starting point is the de-generated  
 178 hyper-rectangle composed of the single point given by the parameter values of the original  
 179 plan. The algorithm tries to extend the hyper-rectangle along one dimension (i.e. it tries  
 180 to widen the interval of possibilities associated to one of the parameters) and checks if the  
 181 resulting hyper-rectangle is in fact an under-approximation of the RE. If it is, the new  
 182 hyper-rectangle is kept as it is guaranteed to be a valid DRE. Otherwise, another dimension  
 183 or another increment is chosen for the algorithm to proceed. The general intuition behind  
 184 the algorithm is depicted in Figure 2.

185 Algorithm 1 reports the pseudo-code of IRR. The formula  $enc_{valid}$  is computed once and  
 186 off-line. It corresponds to the basic requirements for the hyper-rectangle to be a valid DRE:  
 187 only parameter values that are not contradicting the STN plan and the causal flow of effects  
 188 are admissible. This is the same as the first piece of the logical formulation in [2], but luckily  
 189 it is the easier part of the quantification and can be efficiently computed. Then, the IRR  
 190 function is in charge of computing a hyper-rectangle  $R$  maintaining the following invariant:  
 191 at each step,  $R$  is a subset of the RE of the problem. The hyper-rectangle  $R$  is represented  
 192 as a pair of bounds (lower- and upper-) assigned to each parameter (this directly models a  
 193 DRE as per Definition 3), and is initialized (line 3) with the values of the non-parametric  
 194 plan  $\pi$ . The algorithm uses two functions to control how the hyper-rectangle is transformed  
 195 from one cycle to the next.  $\Delta$  associates to each parameter a number that is the value used  
 196 to increase the upper-bound or to decrease the lower-bound for that parameter. The initial  
 197 value for  $\Delta$  is the original value of the parameter scaled by a weight for such parameter, but  
 198 any positive number bigger than  $\beta$  is enough to guarantee soundness and termination of the  
 199 algorithm. Note that these weights can be used to express preferences on the parameters: a  
 200 higher weight pushes the algorithm to expand a specific parameter more than others. The  
 201 function  $\Theta$  is used to decide in which direction the interval of a parameter can be extended.

202 Two directions are possible: UB indicates that we want to extend the upper-bound and LB  
 203 indicates that we want to decrease the lower-bound (line 10). Initially both directions are  
 204 possible, but when we discover (line 13) that one direction is infeasible with the current  
 205  $\Delta$ , we remove this direction from the possibilities. This process will be continued with the  
 206 current  $\Delta$  values until expansion in all directions is infeasible. At this stage the value of  
 207  $\Delta$  is halved, eventually reaching a value lower than  $\beta$ . Each time that the values of  $\Delta$  are  
 208 updated, we reset  $\Theta$  to allow expansion in both directions once again.

209 The main loop of the algorithm continues until all the values of  $\Delta$  are lower than  $\beta$ :  
 210 this is to guarantee that the minimum distance from each border of the hyper-rectangle  
 211 and the border of the actual RE is at most  $\beta$ . The algorithm picks a parameter  $\tilde{\gamma}$  to be  
 212 analyzed among the parameters having at least one direction available in  $\Theta$  and that have not  
 213 converged already (line 7); then, it generates a candidate hyper-rectangle  $R'$  by extending  
 214 either the lower- or the upper- bound of  $\tilde{\gamma}$ . At this point, we check if  $R'$  is contained in the  
 215 RE or not (line 12). If it is, we keep it and continue the loop, otherwise, we discard this  
 216 hyper-rectangle and we record that with the current  $\Delta$  we cannot extend  $\tilde{\gamma}$  in this direction  
 217 by removing the direction  $\theta$  from  $\Theta(\tilde{\gamma})$ . Moreover, if no direction is left for  $\tilde{\gamma}$ , we halve its  
 218 value of  $\Delta$  and reset  $\Theta$  so that  $\tilde{\gamma}$  can be tentatively extended again using a smaller step (lines  
 219 15-16).

220 The core part of the algorithm consists in checking a candidate hyper-rectangle for  
 221 containment in the actual RE, without explicitly computing the region itself. This is done  
 222 via the CHECKINENVELOPE function that performs two SMT checks corresponding to the  
 223 two quantifiers appearing in the RE logical formulation of [2]. The first check looks for points  
 224 belonging to  $R$  that are not parts of the validity region  $enc_{valid}$ , the second checks if the  
 225 rectangle (together with the guarantees from the plan and the effects) implies the proof  
 226 requirements characterizing the REs. The important point here, is that both checks are  
 227 quantifier-free, i.e. no quantifier elimination is involved.

228 ► **Theorem 4.** *The CHECKINENVELOPE( $R$ ) function returns true if and only if  $R$  is a valid*  
 229 *DRE*

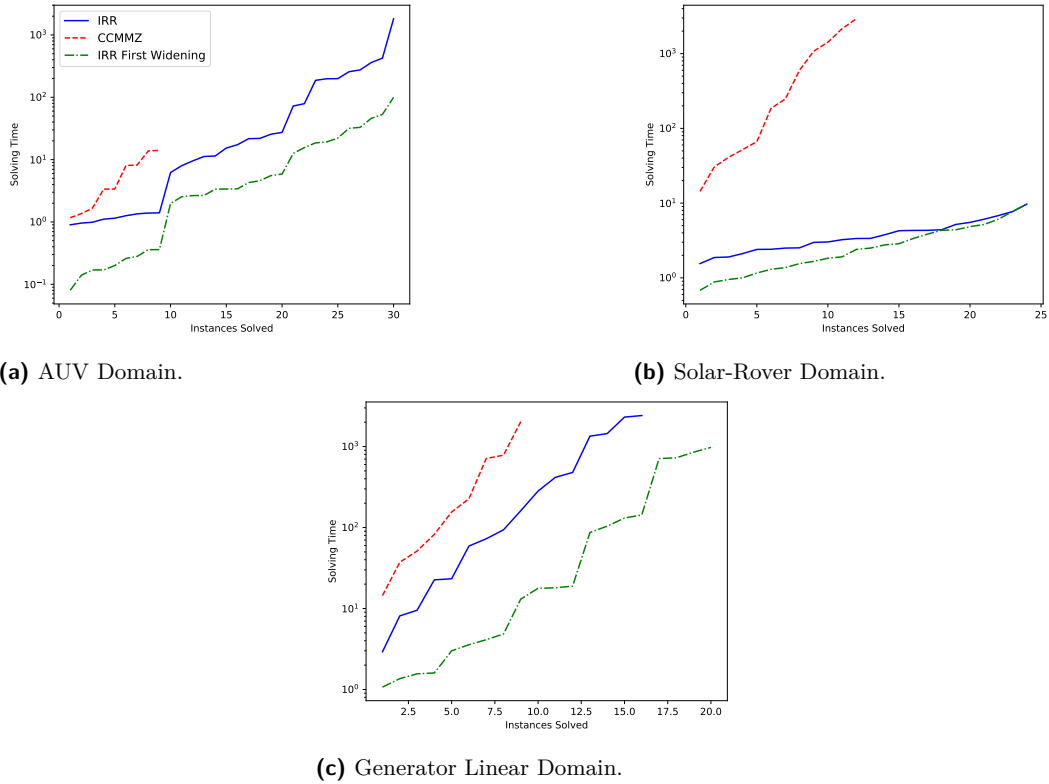
230 **Proof.** The algorithm logically checks the following formula:  $\neg(\exists \bar{\Gamma}. enc_R \wedge \neg enc_{valid}) \wedge$   
 231  $\forall \bar{\Gamma}, \bar{X}. (enc_{tn}^{\pi r} \wedge enc_{eff}^{\pi r} \wedge enc_R) \rightarrow enc_{proofs}^{\pi r}$ , that can be rewritten as  $\forall \bar{\Gamma}. enc_R \rightarrow (enc_{valid} \wedge$   
 232  $(\forall \bar{X}. (enc_{tn}^{\pi r} \wedge enc_{eff}^{\pi r}) \rightarrow enc_{proofs}))$  that states that  $enc_R$  is a subset of the encoding of the  
 233 RE. Then, for Definition 3,  $R$  is the encoding of a valid DRE. ◀

234 An interesting feature of the algorithm is that it is “anytime”, i.e. at each time, we can  
 235 take the hyper-rectangle  $R$  and we have the guarantee that  $R$  is contained in the RE and is  
 236 thus a valid DRE. Moreover, the algorithm is guaranteed to terminate if the RE is finite in  
 237 all dimensions.

238 ► **Theorem 5.** *If the robustness envelope is bounded in all dimensions, IRR always termin-*  
 239 *ates.*

240 **Proof.** All the values in  $\Delta$  are initially positive and whenever the candidate rectangle is  
 241 found to exit the RE (line 13) one of the values in  $\Delta$  is halved. Eventually all the parameters  
 242 will be considered and they will be eventually found to exit the RE because it is bounded in  
 243 all dimensions. Therefore, all the values of  $\Delta$  will become smaller than  $\beta$ . ◀

244 We highlight that IRR is in fact an optimization procedure that incrementally maximizes  
 245 the size of a starting DRE, terminating when a maximal DRE is found within the given  
 246 precision limit  $\beta$ .



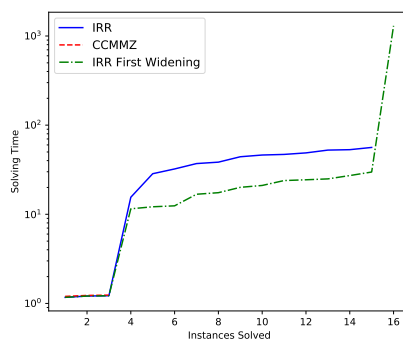
■ **Figure 3** Scalability experiments on [Cashmore *et al.*, 2019] domains: the number of solved instances (sorted by difficulty for each solver) is considered on the X axis and is compared with the logarithmic time needed to solve each instance (lower, longer lines are better).

## 247 5 Experiments

248 We now present our empirical analysis that comprises three sets of experiments. The first  
 249 aims at showing the superior performance of IRR as compared with the logical approach  
 250 of [2]. The second shows a practical use-case of the execution flow proposed in this paper  
 251 when the duration of actions is uncertain. In the third experimentation we use our DRE  
 252 technique to execute plans when the consumption rates of resources is uncertain.

253 **IRR.** We start by considering the experimental dataset and the tool (hereafter called  
 254 CCMMZ) provided in [2]. The benchmarks use a varying number of parameters; in  
 255 particular, AUV ranges from 1 to 8 parameters, Generator Linear from 1 to 4 and Solar  
 256 Rover between 1 and 4. We compare our IRR implementation with CCMMZ on all the  
 257 available instances and domains, measuring the total run-time and using the “decoupled  
 258 envelope generation” functionality of the tool. Moreover, in order to take into account the  
 259 anytime nature of IRR, we also measure the time at which the rectangle  $R$  in IRR widens  
 260 and becomes different than a single point (i.e. we measure the first time the Algorithm 1  
 261 reaches line 17) and we call this timing “IRR First Widening”. In all our experiments we  
 262 set  $\beta = 1$  and all  $\omega_i = 1$  to find the decoupled region approximated to a single unit with no  
 263 preferences among the parameters (obviously, we set the same parameter preference also in  
 264 CCMMZ). We executed all of the instances on a Xeon E5-2620 2.10GHz machine setting a  
 265 time limit of 3600s and a RAM memory limit of 20GB.





■ **Figure 4** Scalability experiments on the delivery domain: the number of solved instances (sorted by difficulty for each solver) is considered on the X axis and is compared with the logarithmic time needed to solve each instance (lower, longer lines are better).

266 Figure 3 shows the result of this analysis. IRR is able to solve many more instances than  
 267 CCMMZ and is consistently quicker. Moreover, we note how the first widening is often  
 268 encountered quite early in the execution, marking the margin for anytime exploitation of  
 269 IRR. In fact, after the first widening, IRR already computed a meaningful and non-trivial  
 270 under-approximation of the RE that can be used for execution. This is particularly evident  
 271 in the Generator Linear domain where the algorithm is unable to fully terminate in some  
 272 cases, but the first widening point is reached.


273 In addition to these domains, we also experiment with several instances of a service-robot  
 274 domain that we also use for the following execution experimental analysis. The domain,  
 275 called “Robot Delivery” is a simplified version<sup>2</sup> of the domain used in the Planning and  
 276 Execution Competition for Logistics Robots in Simulation [16]. The domain comprises a fleet  
 277 of small robots that can navigate in an euclidean graph. These robots are tasked to pick and  
 278 deliver orders within a deadline. Collecting orders requires two robots present at a machine.  
 279 We scaled the number of parameters in the instances between 1 and 33. Figure 4 shows  
 280 the scalability of IRR and CCMMZ on this domain. These instances are much harder for  
 281 both the solvers compared to the previous domains; in fact, CCMMZ is only able to solve 3  
 282 instances, while IRR is able to solve 15 of them. Also in this case, the anytime nature of IRR  
 283 is evident by observing the difference from the first widening and the algorithm completion.

Finally, we investigate how quickly the IRR algorithm converges in our experiments. We  
 define convergence at step  $i$  in a run of IRR that terminates with hyper-rectangle  $R_{end}$  as  
 follows ( $R_i$  indicates the hyper-rectangle at step  $i$ ).

$$Convergence(i) = \frac{\sum_{[l,u] \in R_i} u - l}{\sum_{[l,u] \in R_{end}} u - l} \times 100$$

284 Intuitively, this gives the percentage of the region covered by  $R_i$  with respect to  $R_{end}$ . (Note  
 285 that  $R_{end}$  contains  $R_i$  because the IRR algorithm only expands previous hyper-rectangles.)  
 286 Figure 5 shows, for all problems solved by IRR in our benchmark set, the percentage of  
 287 convergence achieved after any number of steps of the IRR algorithm. Clearly from the plot,

<sup>2</sup> A simplified RCLL domain was used because the PDDL provided in the RCLL image is not complete and the RCLL simulation requires external processes, e.g. a referee box. We are interested in the flexible execution success rate, so we created PDDL instances encoding logistics problems without any external processes.



■ **Figure 5** Convergence of IRR in terms of steps: each red dot is a DRR computed by IRR and the plot shows the progression in terms of convergence at each step of the algorithm. The purple line indicates the poorest convergence percentage for each step in any experiment; similarly the black, dashed line shows the best convergence.

Executor	1 Parameter		2 Parameters		3 Parameters		4 Parameters		5 Parameters	
	Coverage	Avg Replans	Coverage	Avg Replans	Coverage	Avg Replans	Coverage	Avg Replans	Coverage	Avg Replans
DREEx	<b>92.2%</b>	<b>0.1</b>	<b>85.8%</b>	<b>0.2</b>	<b>83.2%</b>	<b>0.2</b>	<b>72.8%</b>	<b>0.1</b>	<b>63.0%</b>	<b>0.1</b>
BLEx(0)	0.0%	NA	0.0%	NA	0.0%	NA	0.0%	NA	0.0%	NA
BLEx(10)	24.5%	1.0	4.8%	1.0	0.8%	1.2	0.2%	0.8	0.0%	NA
BLEx(20)	44.4%	0.8	19.9%	1.0	6.3%	1.0	2.4%	1.6	0.8%	1.9
BLEx(30)	58.9%	0.7	34.0%	0.9	18.8%	1.2	9.2%	1.2	4.5%	1.4
BLEx(40)	68.8%	0.6	52.2%	0.8	34.7%	1.0	23.8%	1.1	11.8%	1.5
BLEx(50)	75.0%	0.5	62.2%	0.7	49.0%	0.9	37.8%	1.1	27.9%	1.2
BLEx(60)	78.8%	0.4	69.0%	0.5	59.4%	0.7	53.4%	0.9	44.5%	1.1

■ **Table 1** Coverage and average number of re-plans in the duration-uncertain delivery domain.

288 in a limited number of steps we often approximate very well the final intervals; in particular,  
 289 within 50 steps we already cover more than 70% of the final sum of the interval sizes in all  
 290 the cases.

291 **Duration-Uncertain Flexible Execution.** We use the Robot Delivery domain to  
 292 investigate the merits of an on-line plan executor equipped with our IRR algorithm. In  
 293 particular, we begin by focusing the analysis on the number of re-plans and on the plan  
 294 execution success rate when only the duration of actions is uncertain during execution. In this  
 295 domain, a robot has to collect a spindle from a shelf, construct a base by performing six steps  
 296 (possibly in parallel), then mount a number of rings, and finally deliver the order. Orders  
 297 have deadlines that must be met for delivery. The domain allows the agent to drop an order  
 298 and restart from scratch with a new one at any time, but this disposal action takes some time  
 299 (10 seconds in our case) and the robot needs to navigate on a symbolic euclidean graph to  
 300 pick the parts, assemble and deliver the order. Each instance is simulated in an environment  
 301 where actions have a non-deterministic duration described by a normal distribution with  
 302 a minimum value. Due to the difficulty in manipulation tasks, the actions executed for  
 303 preparing the base (in which the robots interact with machines) have the highest degree of  
 304 variance. These actions have mean durations of 120, 130, 140, 150, 160 and 170 seconds, and  
 305 a standard deviation of 70. Due to this uncertainty and the presence of deadlines for the  
 306 order delivery, the execution of a plan can fail even when a re-planning schema is employed.  
 307 We generated a total of 100 problems by varying the deadlines for the orders.

---

**Algorithm 2** STN Dispatch
 

---

```

1: function STNDISPATCH( $\pi_{stn}, \rho$ )
2:   finished = false
3:   while  $\neg$ finished do
4:     for each node  $n \in \pi_{stn}$  do
5:        $min, max \leftarrow$  MINMAXDISPATCHTIME( $n, \pi_{stn}, \rho$ )
6:       if  $n$  is action start then
7:         if  $(min \leq n \leq max) \wedge \neg$ STARTED( $n$ ) then
8:           STARTEXECUTING( $n$ )
9:         else if  $(n \geq max) \wedge \neg$ STARTED( $n$ ) then
10:          finished = true
11:        else if  $n$  is action end then
12:          if  $(n \geq max) \wedge \neg$ COMPLETED( $n$ ) then
13:            finished = true
14:          else if  $(n \leq min) \wedge$  COMPLETED( $n$ ) then
15:            finished = true
16:   return GOALSACHIEVED( )

```

---

308 Our DRE-based approach was implemented in ROSPlan, as described in Section 3. The  
 309 STN dispatcher starts the execution of actions following the temporal constraints of the  
 310 STN: the process is illustrated in Algorithm 2. For each node, the minimum and maximum  
 311 dispatch times are calculated during execution (line 5). The dispatch ends when an action  
 312 completes outside of the temporal constraints allowed by the STN, or has not been started  
 313 after the maximum allowed dispatch time. When the dispatch ends, it returns *true* if the  
 314 goals have been achieved; otherwise, re-planning is triggered as shown in Figure 1. The  
 315 system will continuously attempt to re-plan until the deadlines make the PDDL planning  
 316 problem unsolvable.

317 We compare the executor described in Section 3 (indicated as DREEX) against several  
 318 baselines in which we dispatch the STN plan  $\pi_{stn}$  without parameterization. In such baselines,  
 319 the executor dispatches the STN plan allowing for a fixed deviation in the duration of actions  
 320 and ends dispatch only when the action duration falls outside of this interval. This is the  
 321 optimistic technique for execution implemented in ROSPlan that, differently from DREEX,  
 322 offers no formal guarantees. We consider baseline executors named BLEX(0) to BLEX(60)  
 323 allowing for 0% to 60% variability in action duration before triggering a re-plan. For example,  
 324 given an action with a predicted duration 100 seconds, BLEX(0) will re-plan if the duration is  
 325 not exactly 100; BLEX(20) will re-plan if the duration is outside of the interval [80, 120]. The  
 326 baseline BLEX(0) corresponds to formally executing the time-triggered plan  $\pi_{tt}$ : re-planning  
 327 happens if any action duration differs from what was expected in  $\pi_{tt}$ . We highlight that,  
 328 when DREEX is employed and the observation is within the envelope computed by IRR,  
 329 we have the formal guarantee of plan success; as soon as one observation is outside of the  
 330 envelope, we choose to re-plan.

331 The overarching idea in these experiments is that the planner usually optimistically  
 332 selects the easier, quicker goal and the agent starts to execute the plan. If the execution of  
 333 the preparation actions goes overlong, it might become impossible to deliver the order, so  
 334 the only way to successfully recover is to immediately dispose the current order and switch  
 335 to another one with a less imminent deadline. If the executor fails in realizing this situation,  
 336 it continues to execute the plan until it tries to deliver the order, at which point it realizes  
 337 that the deadline is not met. Since a lot of time has been wasted in the preparation, it might  
 338 be impossible to recover from this situation. Ideally, we expect that the predictive power of

Executor	Coverage	Avg Replans
DREEX	<b>99.2%</b>	<b>0.1</b>
BLEX(0)	1.0%	2.0
BLEX(10)	25.6%	<b>0.1</b>
BLEX(20)	50.7%	<b>0.1</b>
BLEX(30)	66.4%	<b>0.1</b>
BLEX(40)	77.9%	<b>0.1</b>
BLEX(50)	82.1%	<b>0.1</b>
BLEX(60)	86.8%	<b>0.1</b>

■ **Table 2** Coverage and average number of re-plans in the resource-uncertain delivery domain.

339 DREs allows the identification of situations where the deadline cannot be met and a swift  
 340 re-planning to change the objective order is needed.

341 Table 1 reports the results of the experiment. We report the coverage percentage (i.e.  
 342 the percentage of problems successfully executed over the benchmark set) as well as the  
 343 average number of re-plannings for successful runs. The baseline BLEX(0), not accounting  
 344 for any variance in action duration, was unable to solve any problem successfully. Allowing  
 345 for more flexibility in the duration of actions increases the coverage as should be expected.  
 346 However, the DREEX approach achieves greater coverage than all baselines in all the cases.  
 347 This is because in this problem, the ability to realize early that the agent is late for the first  
 348 order and change course of actions to achieve the second order is pivotal for achieving a good  
 349 success rate.

350 **Resource-Uncertain Flexible Execution.** Finally, we show that our flow can be used  
 351 when parameters are not just action durations. We expanded the delivery domain to consider  
 352 the battery consumption of the robots. In particular, each action in the revised domain  
 353 checks that enough battery is present upon start and consumes a fixed amount of battery.  
 354 We parametrized the consumption rate of actions, so that the DRE will compute the possible  
 355 consumption values for which a given plan is valid. The executor is then demanded to observe  
 356 the contingent consumption and possibly invoke a re-planning if the observation does not  
 357 fall in the DRE prescription. Also in this case, the baselines BLEX(X) invoke the replanning  
 358 when the battery consumption is observed to be X% higher or lower than the nominal value.

359 Table 2 reports the results of the experiment, and shows how the use of DREEX  
 360 is beneficial for the success-rate achieving an almost perfect success-rate with very few  
 361 replannings on average.

## 362 6 Conclusion

363 In this paper, we make the case for the use of Robustness Envelopes (RE) in a plan execution  
 364 framework. We present a novel, anytime algorithm to compute Decoupled Robustness  
 365 Envelopes (DRE) that is empirically superior to the previously known logic-based construction.  
 366 Moreover, we demonstrate the usefulness of the produced artifacts by integrating them in the  
 367 ROSPlan framework and showing on a concrete example the positive impact on the number  
 368 of re-plannings and the plan success-rate.

369 In the future, we will consider other kinds of approximations for the robustness envelope  
 370 (e.g. hyper-octagons instead of hyper-rectangles). We will also explore the link to temporal  
 371 uncontrollability and non-deterministic planning. Finally, using Incremental Rectangular-  
 372 Robustification (IRR) in parallel with the dispatcher could allow variation in parameters  
 373 being considered during execution.

374 ——— **References** ———

- 375 1 C. Barrett, R. Sebastiani, S. Seshia, and C. Tinelli. Satisfiability modulo theories. In *Handbook*  
376 *of Satisfiability*, pages 825–885. IOS Press, 2009.
- 377 2 M. Cashmore, A. Cimatti, D. Magazzeni, A. Micheli, and P.a Zehtabi. Robustness envelopes  
378 for temporal plans. In *AAAI*, 2019.
- 379 3 M. Cashmore, M. Fox, D.Long, D. Magazzeni, B. Ridder, A. Carrera, N. Palomeras, N. Hurtós,  
380 and M. Carreras. Rosplan: Planning in the robot operating system. In *Proceedings of the*  
381 *Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015,*  
382 *Jerusalem, Israel, June 7-11, 2015.*, pages 333–341, 2015.
- 383 4 A. Cesta, G. Cortellessa, S. Fratini, A. Oddi, and R. Rasconi. The APSI Framework: a Planning  
384 and Scheduling Software Development Environment. In *ICAPS (Application Showcase)*, 2009.
- 385 5 A. Cimatti, M. Do, A. Micheli, M. Roveri, and D. Smith. Strong temporal planning with  
386 uncontrollable durations. *Artif. Intell.*, 256:1–34, 2018. doi:10.1016/j.artint.2017.11.006.
- 387 6 R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*,  
388 49(1-3):61–95, 1991.
- 389 7 M. Do and S. Kambhampati. Improving temporal flexibility of position constrained metric  
390 temporal plans. In *ICAPS*, pages 42–51, 2003.
- 391 8 M. Fox and D. Long. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning  
392 Domains. *Journal of Artificial Intelligence Research*, 20:61–124, 2003.
- 393 9 J. Frank and A. Jónsson. Constraint-based Attribute and Interval Planning. *Constraints*,  
394 8(4):339–364, 2003.
- 395 10 J. Frank and P. Morris. Bounding the resource availability of activities with linear resource  
396 impact. In *ICAPS*, pages 136–143, 2007.
- 397 11 M. Ghallab and H. Laruelle. Representation and control in IxTeT, a temporal planner. In  
398 *AIPS*, pages 61–67, 1994.
- 399 12 M. Ghallab, D. Nau, and P. Traverso. *Automated planning - theory and practice*. Elsevier,  
400 2004.
- 401 13 Félix Ingrand and Malik Ghallab. Deliberation for autonomous robots: A survey. *Artificial*  
402 *Intelligence*, 247:10–44, 2017.
- 403 14 Mausam and A. Kolobov. Planning with markov decision processes: An ai perspective.  
404 *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–210, 2012. doi:  
405 10.2200/S00426ED1V01Y201206AIM017.
- 406 15 N. Muscettola. Computing the envelope for stepwise-constant resource allocations. In *CP*,  
407 pages 139–154, 2002.
- 408 16 T. Niemueller, G. Lakemeyer, and A. Ferrein. The robocup logistics league as a benchmark  
409 for planning in robotics. *Planning and Robotics (PlanRob-15)*, page 63, 2015.
- 410 17 N. Policella, S. Smith, A. Cesta, and A. Oddi. Generating robust schedules through temporal  
411 flexibility. In *ICAPS*, pages 209–218, 2004.
- 412 18 A. Schrijver. *Theory of Linear and Integer Programming*. J. Wiley & Sons, 1998.
- 413 19 A. Umbrico, A. Cesta, M. Cialdea Mayer, and A. Orlandini. Integrating resource management  
414 and timeline-based planning. In *ICAPS*, pages 264–272, 2018.
- 415 20 B. Williams and V. Gupta. Unifying model-based and reactive programming within a model-  
416 based executive. In *Workshop on Principles of Diagnosis*, 1999.

417 **A Planning Semantics**

418 In this section we reproduce the additional syntax and semantics of planning problems  
419 presented in [2], which provide additional background on definitions 1 and 2. For a more  
420 complete description linking these concepts we refer the reader to the source. We start by  
421 defining our planning language: we adopt the full PDDL 2.1 [8] with continuous change.

422 ► **Definition 6.** A *planning problem*  $\mathcal{P}$  is a tuple  $\langle P, V, A, I, G \rangle$ , where  $P$  is a set of  
 423 propositions;  $V$  is a set of real variables, called *fluents*;  $A$  is a set of durative and instantaneous  
 424 actions;  $I : P \cup V \rightarrow \{\top, \perp\} \cup \mathbb{R}$  is the total function describing the initial state of the predicates  
 425 and the fluents.  $G : P \cup V \rightarrow \{\top, \perp\} \cup \mathbb{R}$  is a (possibly partial) function indicating the goal  
 426 condition. A durative action  $a$  is a tuple  $\langle pre_a, eff_a, dur_a \rangle$ , where  $pre_a$  is a set of conditions  
 427 for the actions partitioned in three subsets  $pre_{\vdash a}$ ,  $pre_{\leftrightarrow a}$  and  $pre_{\dashv a}$  of at-start, over-all and  
 428 at-end conditions;  $eff_a$  is the set of action effects, partitioned in seven sets:  $eff_{\vdash a}^+$  (positive  
 429 starting effects),  $eff_{\vdash a}^-$  (negative starting effects),  $eff_{\vdash a}^{num}$  (numeric starting effects),  $eff_{\dashv a}^+$   
 430 (positive ending effects),  $eff_{\dashv a}^-$  (negative ending effects),  $eff_{\dashv a}^{num}$  (numeric ending effects)  
 431 and  $eff_{\leftrightarrow a}^{num}$  (continuous numeric effects); and  $dur_a$  is a set of duration constraints. An  
 432 instantaneous action  $a$  is a tuple  $\langle pre_a, eff_a \rangle$ , where  $pre_a$  is a set of pre-conditions and  $eff_a$   
 433 is the set of action effects, partitioned in  $eff_a^+$  (positive effects),  $eff_a^-$  (negative effects) and  
 434  $eff_a^{num}$  (numeric effects).

435 In the usual PDDL 2.1 setting, a plan is defined as a set of actions associated with a starting  
 436 time and a duration. We define this kind of plans as time-triggered plans.

437 ► **Definition 7.** A *time-triggered plan*  $\pi$  for a planning problem  $\mathcal{P} \doteq \langle P, V, A, I, G \rangle$  is a  
 438 set of tuples  $\langle t, a, d \rangle$ , with  $t \in \mathbb{R}_{\geq 0}$ ,  $a \in A$  and  $d \in \mathbb{R}_{> 0}$  iff  $a$  is a durative action.

439 For the sake of brevity, we omit the formal definition of validity for such a plan, which can  
 440 be found in [8]. Here, it suffices to remind oneself that a plan is valid if by simulating the  
 441 system controlled by the plan, all the prescribed actions are applicable (all their conditions  
 442 are satisfied at the time the action is executed) and the goal is reached after the last action  
 443 terminates.

444 We define an STN plan as a constraint network of time points indicating the starting or  
 445 the ending of actions. Note that the STN plan contains all the information of, and is strictly  
 446 more general than a time-triggered plan. Moreover, that it is not necessary to first find a  
 447 time-triggered plan in order to generate an STN plan.

448 ► **Definition 8.** An *STN plan*  $\pi$  for  $\mathcal{P} \doteq \langle P, V, A, I, G \rangle$  is a tuple  $\langle T, C \rangle$ , where  $T$  is the set  
 449 of time points  $\{z\} \cup \{t_{da}^s, t_{da}^e \mid da \text{ is a durative action instance}\} \cup \{t_a \mid a \text{ is an instantaneous}$   
 450  $\text{action instance}\}$  and  $C$  is a set of constraints in the form  $t_i - t_j \leq b$  with  $t_i, t_j \in T$ , and  
 451  $b \in \mathbb{R}$ .

452 Finally, we can define the validity of an STN plan by considering the set of all possible  
 453 time-triggered plans that are compatible with the STN specification. If all such plans are  
 454 valid, we say that the STN plan is valid.

455 ► **Definition 9.** Given an STN plan  $\pi \doteq \langle T, C \rangle$  and an assignment  $\mu : T \rightarrow \mathbb{R}$  s.t.  
 456  $\mu(z) = 0$ , the *induced time-triggered plan* by  $\mu$  is the time-triggered plan  $tt(\mu) \doteq$   
 457  $\{\langle \mu(t_{da}^s), da, \mu(t_{da}^e) - \mu(t_{da}^s) \rangle \mid da \text{ is a durative action}\} \cup \{\langle \mu(t_a), a, 0 \rangle \mid a \text{ is an instantaneous}$   
 458  $\text{action}\}$ .

459 ► **Definition 10.** An STN plan  $\pi \doteq \langle T, C \rangle$  for  $\mathcal{P}$  is *valid* if for each assignment  $\mu : T \rightarrow \mathbb{R}$   
 460 s.t.  $\mu(z) = 0$  and for all  $t_i - t_j \leq b \in C$   $\mu(t_i) - \mu(t_j) \leq b$ , the time-triggered plan  $tt(\mu)$  is  
 461 valid for  $\mathcal{P}$ .