



Heriot-Watt University
Research Gateway

Proof-Carrying Plans: a Resource Logic for AI Planning

Citation for published version:

Hill, A, Komendantskaya, E & Petrick, RPA 2020, Proof-Carrying Plans: a Resource Logic for AI Planning. in *PPDP '20: Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming.*, 14, Association for Computing Machinery, 22nd International Symposium on Principles and Practice of Declarative Programming 2020, Bologna, Italy, 8/09/20.
<https://doi.org/10.1145/3414080.3414094>

Digital Object Identifier (DOI):

[10.1145/3414080.3414094](https://doi.org/10.1145/3414080.3414094)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Peer reviewed version

Published In:

PPDP '20: Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming

Publisher Rights Statement:

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Proof-Carrying Plans: a Resource Logic for AI Planning

Alasdair Hill
Heriot-Watt University
Edinburgh, Scotland
ath7@hw.ac.uk

Ekaterina Komendantskaya
Heriot-Watt University
Edinburgh, Scotland
ek19@hw.ac.uk

Ronald P. A. Petrick
Heriot-Watt University
Edinburgh, Scotland
R.Petrick@hw.ac.uk

Abstract

Planning languages have been used successfully in AI for several decades. Recent trends in AI verification and Explainable AI have raised the question of whether AI planning techniques can be verified. In this paper, we present a novel resource logic, the *Proof Carrying Plans (PCP) logic* that can be used to verify plans produced by AI planners. The PCP logic takes inspiration from existing resource logics (such as Linear logic and Separation logic) as well as Hoare logic when it comes to modelling states and resource-aware plan execution. It also capitalises on the Curry-Howard approach to logics, in its treatment of plans as functions and plan pre- and post-conditions as types. This paper presents two main results. From the theoretical perspective, we show that the PCP logic is sound relative to the standard possible world semantics used in AI planning. From the practical perspective, we present a complete Agda formalisation of the PCP logic and of its soundness proof. Moreover, we showcase the Curry-Howard, or functional, value of this implementation by supplementing it with the library that parses AI plans into Agda's proofs automatically. We provide evaluation of this library and the resulting Agda functions.

Keywords: AI planning, Verification, Resource Logics, Theorem Proving, Dependent Types.

CCS Concepts

• **Theory of computation** → **Action semantics; Operational semantics; Logic and verification;** • **Computing methodologies** → **Planning for deterministic actions;** • **Software and its engineering** → **Formal software verification.**

ACM Reference Format:

Alasdair Hill, Ekaterina Komendantskaya, and Ronald P. A. Petrick . 2020. Proof-Carrying Plans: a Resource Logic for AI Planning. In *22nd International Symposium on Principles and Practice of Declarative Programming (PPDP '20)*, September 8–10, 2020, Bologna, Italy. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3414080.3414094>

1 Motivation

Planning is a research area within AI that studies automated generation of plans from symbolic domain and problem specifications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PPDP '20, September 8–10, 2020, Bologna, Italy

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8821-4/20/09...\$15.00

<https://doi.org/10.1145/3414080.3414094>

```
(define (problem blocksworld)
  (:domain blocksworld)
  (:objects a b)
  (:init (onTable a)f
         (onTable b)
         (clear a)
         (clear b)
         (handEmpty))
  (:goal (and (on a b) (onTable b))))
```

Figure 1: “BlocksWorld” Planning Problem Description. Initial state: two blocks, a and b, are lying “clear” (i.e. unobstructed) on a table, and a robot hand is empty. Goal state: block a is on block b.

AI planners came into existence in the 1970s as an intersection between general problem solvers [11], situation calculus [28] and theorem proving [17]. One of the most popular early planners was the Stanford Research Institute Problem Solver (STRIPS) [12] which was created to address the problems faced by a robot in rearranging objects and in navigating.

In STRIPS, a planner is given a description of an *initial state* (of the “world”) and a *goal state*. For example, Figure 1 defines the initial state that has blocks a and b on the table, and the goal state – the blocks assembled in a stack. A solution to a planning problem is a sequence of actions, which is simply referred to as a *plan*. For example, a solution to the planning problem of Figure 1 is the following plan: *pickup a from the table*, then *putdown a on b*.

Many versions of planning languages were proposed, and the *Planning Domain and Definition Language (PDDL)* [29] aimed to standardise them. One notable design decision of PDDL is the splitting of the planning problem into *domain* and *problem descriptions*. The domain description describes generally predicates and admissible actions (as shown in Figure 2), while the problem description defines specific initial and goal states (as shown in Figure 1).

PDDL has many extensions over regular STRIPS syntax with the latest version supporting types, numerical functions, equality, conditionals, concurrency, temporal planning and more. Among applications are: reasoning about knowledge, belief and causality, planning allocation of resources, modelling perception of the real world, program synthesis and implementations of multi-agent systems [19, 48]. Many of these applications of planning are used in real-world environments where the verification of plan correctness is essential for successful and safe operation.

Verification and validation of AI planning languages [4] is a rich field of research. One may verify domain models, planning algorithms, or the produced plans.

```

(define (domain blocksworld)
  (:requirements :strips :equality)
  (:predicates
    (handEmpty)
    (holding ?x)
    (onTable ?x)
    (on ?x ?y)
    (clear ?x))

  (:action pickup_from_table
    :parameters
      (?x)
    :precondition
      (and (handEmpty)
            (onTable ?x)
            (clear ?x))
    :effect
      (and (not (handEmpty))
            (not (onTable ?x))
            (holding ?x)))

  (:action putdown_on_stack
    :parameters
      (?x ?y)
    :precondition
      (and
        (not (= ?x ?y))
        (holding ?x)
        (clear ?y))
    :effect
      (and
        (not (holding ?x))
        (not (clear ?y))
        (on ?x ?y)
        (handEmpty))))

```

Figure 2: A fragment of PDDL “BlocksWorld” Domain.

Verifying domain models [27, 38] seeks to validate whether domain descriptions accurately capture (expert) knowledge about the world. This can be done by performing test based verification of input and output specifications to check the domain performs as expected. Alternatively, some approaches ensure properties that should hold across many domains such as enforcing that the actions cannot lead to an inconsistent state.

Formalisation of planning algorithms [1] has shown that even well understood algorithms can produce incorrect plans. Modern AI planners are complex software artefacts, and the existing attempts [37, 43] to verify them only focused on certain aspects of their implementation. Due to the complexity of planning problems, many planners will opt for implementations where efficiency is the primary concern which can further complicate the ability to formally verify these algorithms. No mainstream planner has been fully verified yet.

AI plan verification seeks to verify plans produced by planners against some domain model. These tools check properties such as precondition satisfaction, termination and goal satisfaction to ensure that a plan is valid. For example, PDDL has a validator [23] that performs these checks and suggests repairs. This more practical and *lightweight* approach to verification is broadly in line with other lightweight verification trends in the literature [13]. However, at the same time, it is rather disjoint from the growing body of research into *type-based verification* [26, 32, 34] or *resource logics* [6, 41] that offer more principled, formal and rigorous approaches, as well as richer languages for expressing the verification properties.

In this paper, we are taking an attempt at bridging this gap between the AI planning and the programming language community. We introduce a new formal system inspired by resource semantics [36, 41], and by the Curry-Howard view on Separation logic as given e.g. in [30, 40]. We call the resulting formalism the *proof-carrying plan logic* (or PCP logic for short). It features: Hoare triples to describe plans and states; the frame rule for local resource-aware reasoning; and the Curry-Howard view on states as types, and state transformations as functions. The latter feature ensures that plans that we verify in our logic are also executable functions – which completes the analogy with the “proof-carrying code” research agenda [34].

This approach has several advantages over the existing plan verifiers. Firstly, the clear and intuitive formal semantics helps to clarify the computational properties of AI plans. For example, conditions are embedded into our rules that ensure the desired property of state consistency is inherent in the logic; and structural rules of the PCP logic help to clarify the role of constraints in PDDL and expose some latent properties of AI plans (see Section 4). As a result, we were able to prove soundness of the PCP logic relative to the possible world semantics as used in the AI planning community, and fully formalise both the logic and the proof in Agda [20]. This sets up new standards of rigour for AI plan verification that is not present in existing planners.

Secondly, some benefits arise as a consequence of adopting a higher level of abstraction. For example, the existing AI planning verification approaches split into methods for domain and plan verification. This is potentially harmful, as verifying just one aspect still leaves a gap for bugs and errors. The PCP logic does not separate the problem of state consistency and validity of plan execution. We envisage that the right level of abstraction will enable further extensions to incorporate concurrency and more sophisticated constraints on the states and the plans.

Finally, benefiting from the Curry-Howard approach, our Agda code can be extracted as verified executable Haskell or byte code. We will illustrate all of these concepts by means of an example.

1.1 Results of this paper by means of an example

Figures 1, 2 show the original PDDL syntax for a planning domain and a planning problem. PDDL will be able to automatically find a plan that satisfies pre- and post-conditions shown in Figure 1. In particular, it will find a plan

$$f_{ab} = ((\text{pickup_from_table } a) ; (\text{putdown_on_stack } a \ b)).$$

Our goal is to formulate a proof system in which we can (semi-) automatically prove correctness of this plan, given the PDDL domain description, the initial and the goal states.

Looking closer at the domain definition in Figure 2, we see it declares first-order predicates, as well as actions that operate on pre- and post-conditions. Ignoring temporarily the internal structure of pre- and post-conditions, we can see that the formalism lends itself naturally to the syntax of Hoare triples [21]: $\{\text{Pre}\} \rightsquigarrow \{\text{Post}\} \mid \text{action}$. This is our first key intuition. Somewhat differently from Hoare logic, we see that the domain definition defines a set of axioms that control actions. So, we will be talking about a certain plan or action possible relative to a domain Γ . Thus, we will in fact be working with judgements of the form:

$$\Gamma \vdash \{\text{Pre}\} \rightsquigarrow \{\text{Post}\} \mid \text{action}$$

Let us now look closer at the structure of the pre- and post-conditions. The domain specification (Figure 2) states them as formulae of first-order logic (with negation and conjunction), but the problem definition (Figure 1) uses lists of atomic propositional formulae to describe initial and goal states. A simple way to resolve this mismatch is to define Pre and Post to be states in which each individual atomic formula is mapped to + or -, depending on whether it is considered to be true or false in the state. For example, $\{\text{onTable } a \mapsto +\}$ is a state with one formula map. This allows us to formalise the notions of negation, state and (later) action on states.

Of course, we must not allow inconsistent states where a formula is mapped to + and - simultaneously. To ensure this we introduce a notion of a *valid* state where a state is valid if it is consistent, i.e. contains no conflicting formula maps. For example $\{\{\text{onTable } a \mapsto +\} * \{\text{onTable } b \mapsto +\}\}$ is a valid state but $\{\{\text{onTable } a \mapsto +\} * \{\text{onTable } a \mapsto -\}\}$ is not.

States are not necessarily propositional and in particular Figure 2 implies first-order language in the domain definitions. Therefore, we will assume that all $\{\text{Pre}\} \rightsquigarrow \{\text{Post}\} \mid \text{action}$ in Γ are implicitly universally quantified as follows: $\forall \bar{x}. \{\text{Pre}(\bar{x})\} \rightsquigarrow \{\text{Post}(\bar{x})\} \mid \text{action}(\bar{x})$.

One final caveat exists. If we look closer at the domain description of Figure 2, we will notice that it uses an inequality constraint $x \neq y$ that is not declared as a domain predicate. We will follow the resource logics tradition [5, 40] and separate state descriptions from constraints on states. That is, we further refine domain descriptions to have the syntax $\phi; \{\text{Pre}\} \rightsquigarrow \{\text{Post}\} \mid \text{action}$, where ϕ defines constraints on states. For technical reasons, we formalise ϕ to be a list (rather than a conjunction) of constraints, and we use equality ($=$) and inequality (\neq) constraints instead of using negation explicitly. Figure 3 defines the context Γ_{BW} that matches PDDL domain description of Figure 2 in this new language.

It now remains to formulate the rules for the PCP logic. They are very simple: we need a rule “ApplyAction” to be able to choose specific action definitions from the context, we need a rule that composes the actions, similar to the composition rule of Hoare logic [21], and we need a frame rule [9, 18] to have local reasoning on states. Additionally, the system has two structural rules, weakening and shrink. We show that the rules are sound relative to the possible world semantics of PDDL, and we formalise the PCP logic and the soundness proof in Agda [20].

$$\begin{aligned} & \left[\begin{array}{l} \text{handempty} \mapsto + \\ * \text{onTable } x \mapsto + \\ * \text{clear } x \mapsto + \end{array} \right] \rightsquigarrow \left\{ \begin{array}{l} \text{handEmpty} \mapsto - \\ * \text{onTable } x \mapsto - \\ * \text{holding } x \mapsto + \\ * \text{clear } x \mapsto + \end{array} \right\} \mid \alpha_1 x \\ & \text{where } \alpha_1 \equiv \text{pickup_from_table} \\ & [x \neq y]; \left\{ \begin{array}{l} \text{holding } x \mapsto + \\ * \text{clear } y \mapsto + \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{l} \text{holding } x \mapsto - \\ * \text{clear } y \mapsto - \\ * \text{on } x y \mapsto + \\ * \text{handEmpty} \mapsto + \end{array} \right\} \mid \alpha_2 x y \\ & \text{where } \alpha_2 \equiv \text{putdown_on_stack} \end{aligned}$$

Figure 3: Context Γ_{BW} that defines BlocksWorld PDDL domain of Figure 2.

To make use of this Agda library, we can compile the domain and problem definitions from PDDL directly to Agda code. We then can prove in Agda correctness of the PDDL plans. For example, we can prove that, given $P_{ab} \equiv \{(\text{onTable } a \mapsto +) * (\text{onTable } b \mapsto +) * (\text{clear } a \mapsto +) * (\text{clear } b \mapsto +) * (\text{handEmpty} \mapsto +)\}$ and $Q_{ab} \equiv \{(\text{on } a b \mapsto +) * (\text{onTable } b \mapsto +)\}$ as in Figure 2, $\Gamma_{\text{BW}} \vdash P_{ab} \rightsquigarrow Q_{ab} \mid f_{ab}$, i.e. we can certify that the plan f_{ab} is indeed valid.

Finally, we can take advantage of the Curry-Howard interpretation of $\Gamma_{\text{BW}} \vdash P_{ab} \rightsquigarrow Q_{ab} \mid f_{ab}$, as “function f_{ab} has type $P_{ab} \rightsquigarrow Q_{ab}$ ” and actually execute f_{ab} as a function. We define an *action handler*, an auxiliary function that executes plans on states. It will apply the plan f_{ab} to the initial state P_{ab} to obtain the goal state Q_{ab} as a function output. Moreover, we can extract this code to Haskell or binary files, the latter can be deployed directly on robots, with the advantage of carrying the correctness proof! We show the extracted code for this example and several additional examples in [20].

1.2 The Paper Structure

The paper proceeds as follows. Section 2 introduces the PCP logic, proving formally some basic results concerning the ordering and basic operations on states. Section 3 establishes the soundness of the PCP logic and also defines the notion of action handler. Section 4 describes the implementation [20], evaluates it on several benchmark PDDL domains, and discusses the practical value of using dependent types for implementation of verified plans. Section 5 concludes, and discusses related and future work.

2 The PCP Logic

This section defines the syntax, ordering (subtyping) relation on states, and the rules of the PCP logic.

2.1 Syntax of the PCP logic

We define the PCP syntax in Figure 4.

First-order formulas and constraints. Let R be a set of predicate symbols $\{R, R_1, R_2, \dots\}$ with arities, X be a set of variables $\{x, x_1, x_2, \dots\}$, and C be a set of constants $\{c, c_1, c_2, \dots\}$. Figure 4 defines a term as either a variable or a constant. An *atomic formula* (or *Atom*) is given by a predicate applied to a finite list of terms. For example, the atomic formula $\text{onTable } a$ consists of the predicate

onTable applied to a constant a. This defines the pure first-order part of our logic. We also distinguish two specific kinds of atomic formulae that feature equality and inequality as predicate symbols, we call these *Constraints*.

We will use abbreviation \bar{x} to denote a finite list $\{x_1, \dots, x_n\}$ of arbitrary length. We will write $R(\bar{x})$ if R contains variables \bar{x} . A substitution is a partial map from X to C , and we will use symbols $\{\sigma, \sigma_1, \sigma_2, \dots\}$ to denote ground substitutions. Given an atomic formula $R(\bar{x})$ we write $R(\bar{x})[x_i c_i]$ when we substitute each occurrence of a variable x_i in \bar{x} by a constant c_i . We say the resulting formula is *ground*, i.e. it contains no variables.

Actions and plans. Let A be a set of action names $\{\alpha, \alpha_1, \alpha_2, \dots\}$. Figure 4 defines an action as an action name applied to a list of terms, e.g. pickup_from_table a is an action. A plan is a sequence of actions; shrink is a special constructor that can be used in a plan instead of an action, its use will be made clear later.

States and contexts. Polarities $+$ and $-$ are used to denote absence or presence of certain atomic fact in a world. Given a polarity z , $A \mapsto z$ is a *formula map*. A *state* can be given by an empty state, a formula map or a conjunction of such maps (denoted by $*$). A state $(A \mapsto z * P)$ is *valid* if A does not occur in P and P is a valid state. We will only work with valid states in this paper. A *context* Γ contains descriptions of actions in the form $\phi(\bar{x}); \{P(\bar{x})\} \rightsquigarrow \{Q(\bar{x})\} \mid \alpha \bar{x}$ where $\{P(\bar{x})\} \rightsquigarrow \{Q(\bar{x})\}$ denotes a transformation from a state $P(\bar{x})$ to a state $Q(\bar{x})$, $\alpha \bar{x}$ is an action and $\phi(\bar{x})$ is a constraint list.

REMARK ON NOTATION 1. To simplify our notation, we extend the use of notation “ (\bar{x}) ” from atomic formulae, such as $R(\bar{x})$, to states (e.g. $Q(\bar{x})$), actions (e.g. $\alpha(\bar{x})$) and constraints (e.g. $\phi(\bar{x})$). In all these cases, the presence of \bar{x} signifies the presence of free variables \bar{x} in the states, actions, and constraints, respectively. We will drop \bar{x} and will write just Q , α , and ϕ to emphasise that the state, action or constraint do not contain any variables, i.e. they are ground.

A plan specification is a sequent of the form:

$$\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$$

It states that given a context Γ , f is a plan that gives a provable transformation from (ground) state P to (ground) state Q . In the Curry-Howard interpretation of this logic, we view f as a function that inhabits type $\{P\} \rightsquigarrow \{Q\}$.

In all examples, we use the following shorthand notation:

$$R t \mapsto z * R t_1 \mapsto z \equiv R t, t_1 \mapsto z$$

For example, we will write $(\text{onTable } a, b \mapsto +)$ instead of $(\text{onTable } a \mapsto +) * (\text{onTable } b \mapsto +)$. To emphasise that a formula map binds stronger than $*$, we will put parentheses around formula maps in all examples. But we will omit the parentheses in the formal grammar, to keep the notation simple.

2.2 Subtyping (order on states)

We first recall the subtyping relation and the override operator on states introduced in [45], and then establish some lemmas about these, which will be useful in the later sections. The lemmas have not appeared in [45]. We omit proofs here, but give them in Agda [20].

Figure 5 defines order $<$: over states. Following [45], we call it subtyping to refer to the fact that states can also be seen as types. In

this paper, subtyping serves us when we need to compare states or decide whether they are equal. Two states P and Q are considered equal if $P <: Q$ and $Q <: P$.

EXAMPLE 1 (SUBTYPING). Given: $Q \equiv (\text{onTable } a \mapsto -) * (\text{onTable } b \mapsto +) * (\text{clear } a, b \mapsto +) * (\text{handEmpty } \mapsto -) * (\text{holding } a \mapsto +)$ and $Q' \equiv (\text{onTable } a \mapsto -) * (\text{onTable } b \mapsto +) * (\text{clear } a, b \mapsto +) * (\text{holding } a \mapsto +)$, we have $Q <: Q'$.

Subtyping is both reflexive and transitive, i.e. it is a *pre-order*.

LEMMA 1 (SUBTYPING IS PREORDER). Given states P, Q, S , we have:

- (reflexivity) $P <: P$;
- (transitivity) $P <: Q$ and $Q <: S$ implies $P <: S$.

In later sections, we will also need an override operator on states:

DEFINITION 1 (OVERRIDE OPERATOR [45]).

$$P \sqcup \text{emp} = P$$

$$P \sqcup [A \mapsto z * Q] = [A \mapsto z * P \setminus \{A \mapsto + * A \mapsto -\}] \sqcup Q$$

The override operator adds all formula maps from one state to the other. If a mapping for a formula that is to be added already exists, then that formula is removed before adding the new formula map.

EXAMPLE 2 (OVERRIDE OPERATOR).

$$\begin{aligned} & (\text{handEmpty } \mapsto +) * (\text{onTable } a \mapsto +) * (\text{clear } a \mapsto +) \sqcup \\ & (\text{handEmpty } \mapsto -) * (\text{onTable } a \mapsto -) * (\text{holding } a \mapsto +) \\ & = (\text{holding } a \mapsto +) * (\text{onTable } a \mapsto -) * (\text{handEmpty } \mapsto -) * (\text{clear } a \mapsto +) \end{aligned}$$

We have the following lemmas summarising the properties of the subtyping relation and the override operator.

LEMMA 2 (ORDER OF SUBTYPING). Given an atom A and states P and Q , if $A \notin Q$ and $Q <: P$ then $A \notin P$.

LEMMA 3 (MONOTONICITY OF SUBTYPE EXPANSION). Given states P and Q and a formula map $A \mapsto z$, if $Q <: P$ then $A \mapsto z * Q <: P$.

LEMMA 4 (POST-CONDITION OVERRIDE). $(P \sqcup Q) <: Q$ holds for all states P and Q .

LEMMA 5 (MONOTONICITY OF OVERRIDE). Given a polarity z , an atom A , states P and Q , if $A \notin Q$ then $A \mapsto z \in (A \mapsto z * P) \sqcup Q$.

2.3 Normalisation of Constraint Lists

We will now define a normalisation function for constraint lists. This function takes a list of constraints and recurses through them checking that they are true. If a constraint is not true, \perp is returned; otherwise the empty list case will be reached and \top will be returned. We use $t \equiv t_1$ to denote syntactic equivalence between terms.

DEFINITION 2 (NORMALISATION FUNCTION FOR CONSTRAINTS).

$$\text{norm } [] = \top$$

$$\text{norm } (t = t_1 :: \phi) = \text{if } t \equiv t_1 \text{ then norm } \phi \text{ else } \perp$$

$$\text{norm } (t \neq t_1 :: \phi) = \text{if } t \equiv t_1 \text{ then } \perp \text{ else norm } \phi$$

Term	Term $\ni t, t_1, \dots, t_n$	$:= x \mid c$
Atomic Formulae	Atom $\ni A$	$:= R(t_1, \dots, t_n)$
Constraint	Constraint $\ni e$	$:= t = t_1 \mid t \neq t_1$
Constraint List	CList $\ni \phi, \psi$	$:= [] \mid e :: \phi$
Actions	Act $\ni a$	$:= \alpha(t_1, \dots, t_n)$
Plan	Plan $\ni f, f_1, f_2$	$:= \text{shrink} \mid a \mid f; f_1$
Polarities	Polarity $\ni z$	$:= + \mid -$
State	State $\ni P, Q, S$	$:= \text{emp} \mid A \mapsto z \mid P * Q$
(Planning) Context	$\Gamma \ni \gamma$	$:= \phi(\bar{x}); \{P(\bar{x})\} \rightsquigarrow \{Q(\bar{x})\} \mid \alpha \bar{x}$
Specification	Specification $\ni G$	$:= \Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$

Figure 4: The syntax of PCP logic

$$\text{NilSub} \frac{}{S <: \text{emp}} \quad \text{ASub} \frac{S' <: S \quad A \mapsto z \in S'}{S' <: A \mapsto z * S}$$

Figure 5: Subtyping order on states.

EXAMPLE 3 (NORMALISATION FUNCTION FOR CONSTRAINTS). We have $\text{norm}[a = a, b = b] = \top$ but $\text{norm}[a = a, b = c] = \perp$.

2.4 Rules

Figure 6 gives the rules of the PCP logic. We will discuss and illustrate each rule in order, using our running example. In Figure 2 a PDDL definition of BlocksWorld is defined. An example context Γ_{BW} , inspired by that definition, is given in Figure 3. Assume that this is the context for all below examples.

ApplyAction checks that an action is in the context and then constructs the resultant state given by a ground substitution on that action. For example, the `pickup_from_table` action is included in Γ_{BW} (cf. the first action in Figure 3). Taking $P^a \equiv \{(\text{handempty} \mapsto +) * (\text{onTable } a \mapsto +) * (\text{clear } a \mapsto +)\}$ and $Q^a \equiv \{(\text{handempty} \mapsto -) * (\text{onTable } a \mapsto -) * (\text{holding } a \mapsto +) * (\text{clear } a \mapsto +)\}$, we have

$$\frac{(1) \in \Gamma_{\text{BW}}}{\Gamma_{\text{BW}} \vdash \{P^a\} \rightsquigarrow \{Q^a\} \mid \text{pickup_from_table } a}$$

where (1) refers to the first action in Γ_{BW} .

This is the only rule that allows us to access planning domain definitions. Note also that this is the only rule that checks whether constraints on states are satisfied. This is possible thanks to essentially propositional reasoning implemented in planning, thus it is sufficient to check the constraints only once.

Composition rule says that if we have an entailment $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ we can compose it together with another entailment $\Gamma \vdash \{Q'\} \rightsquigarrow \{R\} \mid f_1$ to produce $\Gamma \vdash \{P\} \rightsquigarrow \{R\} \mid f; f_1$, if $Q <: Q'$.

For this example, we take Q and Q' as in Example 1 (with $Q <: Q'$), and we take P and R as follows:

$$P \equiv (\text{onTable } a, b \mapsto +) * (\text{clear } a, b \mapsto +) * (\text{handempty} \mapsto +) \\ R \equiv (\text{onTable } a \mapsto -) * (\text{onTable } b \mapsto +) * (\text{clear } a \mapsto +) * (\text{clear } b \mapsto -) * (\text{handempty} \mapsto +) * (\text{holding } a \mapsto +)$$

\rightarrow) * (on a b \mapsto +). Abbreviating `putdown_on_stack` a b as α and `pickup_from_table` a as f , we have the following application of the Composition rule:

$$\frac{Q <: Q' \quad \Gamma_{\text{BW}} \vdash \{P\} \rightsquigarrow \{Q\} \mid f \quad \Gamma_{\text{BW}} \vdash \{Q'\} \rightsquigarrow \{R\} \mid \alpha}{\Gamma_{\text{BW}} \vdash \{P\} \rightsquigarrow \{R\} \mid f; \alpha}$$

The **Frame** rule allows the addition of formula maps to both states in an entailment, provided the atom of the formula map does not already have a mapping in either state. Continuing the derivation in one of the previous examples, the following application of the frame rule is possible:

$$\frac{\Gamma_{\text{BW}} \vdash \{P^a\} \rightsquigarrow \{Q^a\} \mid \text{pickup_from_table } a}{\Gamma_{\text{BW}} \vdash \{P^a * (\text{onTable } b \mapsto +)\} \rightsquigarrow \{Q^a * (\text{onTable } b \mapsto +)\} \mid \text{pickup_from_table } a} \text{FRAME}$$

In our system, the frame rule is more restrictive than can be seen in other logics such as Separation logic [5, 42], as it can only be used at an action level but not at a plan level. The following example shows the problem with consistency of derivations, if we apply the frame rule to arbitrary judgements of the form $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$.

EXAMPLE 4 (PROBLEMS WITH THE FRAME RULE FOR COMPLEX PLANS). Imagine we have an action α with the transformation $\{\text{clear } a \mapsto +\} \rightsquigarrow \{(\text{clear } a \mapsto -) * (\text{clear } b \mapsto +)\}$ and another action α' with the transformation $\{\text{clear } a \mapsto -\} \rightsquigarrow \{\text{clear } a \mapsto +\}$ then we can compose these two actions together to generate the entailment: $\Gamma \vdash \{\text{clear } a \mapsto +\} \rightsquigarrow \{\text{clear } a \mapsto +\} \mid \alpha; \alpha'$. We have lost the information `clear b \mapsto +` and if the Frame rule was not bound to single actions we could frame incorrectly in the entailment: $\Gamma \vdash \{(\text{clear } a \mapsto +) * (\text{clear } b \mapsto -)\} \rightsquigarrow \{(\text{clear } a \mapsto +) * (\text{clear } b \mapsto -)\} \mid \alpha; \alpha'$, getting a derivation inconsistent with the action definition.

If we want to apply this rule on judgements involving complex plans instead of single actions, then we would need to ensure that the framed atom is not mapped in any state at any level in the plan derivation. This could be done by amending the restrictions

$$\begin{array}{c}
\text{Frame} \frac{\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid \alpha}{\Gamma \vdash \{P * A \mapsto z\} \rightsquigarrow \{Q * A \mapsto z\} \mid \alpha} \\
\text{Shrink} \frac{Q <: Q' \quad \Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f}{\Gamma \vdash \{P\} \rightsquigarrow \{Q'\} \mid f; \text{shrink}} \\
\text{Composition} \frac{Q <: Q' \quad \Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f \quad \Gamma \vdash \{Q'\} \rightsquigarrow \{R\} \mid f_1}{\Gamma \vdash \{P\} \rightsquigarrow \{R\} \mid f; f_1} \\
\text{Weakening} \frac{P' <: P \quad \Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f}{\Gamma \vdash \{P'\} \rightsquigarrow \{Q\} \mid f} \\
\text{ApplyAction} \frac{\phi(\bar{x}); \{P(\bar{x})\} \rightsquigarrow \{Q(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma}{\Gamma \vdash \{P(\bar{x})[\sigma]\} \rightsquigarrow \{Q(\bar{x})[\sigma]\} \mid \alpha(\bar{x})[\sigma]} \\
\text{Where } \phi(\bar{x})[\sigma] \text{ normalises to } \top
\end{array}$$

Figure 6: Rules of the PCP logic. The rules operate on valid states.

on the frame rule or by amending the other rules to prevent loss of information.

Weakening is applied before composition, when a formula map we want in the precondition P already exists in the (previously obtained) post-condition Q . The above example shows a use case with the action α as defined above:

$$\frac{\Gamma_{\text{BW}} \vdash \{(\text{clear } b \mapsto +) * (\text{holding } a \mapsto +)\} \rightsquigarrow \{(\text{clear } b \mapsto -) * (\text{holding } a \mapsto -) * (\text{on } a \text{ } b \mapsto +) * (\text{handEmpty } \mapsto +)\} \mid \alpha}{\Gamma_{\text{BW}} \vdash \{(\text{clear } b \mapsto +) * (\text{holding } a \mapsto +) * (\text{handEmpty } \mapsto -)\} \rightsquigarrow \{(\text{clear } b \mapsto -) * (\text{holding } a \mapsto -) * (\text{on } a \text{ } b \mapsto +) * (\text{handEmpty } \mapsto +)\} \mid \alpha}$$

In BlocksWorld it is implied that `handEmpty` is false when *holding* any block is true and vice versa. This leads the `putdown_on_stack` action's preconditions to only contain the precondition that holding a block has to be true and we use weakening to gain back the information that `handEmpty` is false.

Shrink allows us to shrink and reorder the post-condition state. Any postcondition state Q can be replaced with Q' as long as it is a subtype of the current post state. Shrink can appear anywhere in a plan but currently the main use of this rule is when we have a goal state that is smaller than the obtained post-condition, for example:

$$\frac{\Gamma \vdash \{P\} \rightsquigarrow \{(\text{clear } b \mapsto -) * (\text{holding } a \mapsto -) * (\text{on } a \text{ } b \mapsto +) * (\text{handEmpty } \mapsto +)\} \mid f}{\Gamma \vdash \{P\} \rightsquigarrow \{(\text{on } a \text{ } b \mapsto +)\} \mid f; \text{shrink}}$$

Frame, Weakening and Shrink are structural rules, i.e. they do not change the computational properties of plans, and with the exception of Shrink, do not change the plans syntactically. We finish this section by stating two lemmas that explain subtyping for plans derived by structural rules. Note that all actions have unique definitions in any given context Γ . The proofs of these lemmas are given in Agda [20].

LEMMA 6 (PROPERTY OF STRUCTURAL RULES (LEFT)). *If there is a derivation for $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid \alpha$ by the rules of Figure 6 we have: $\{P'(\bar{x})\} \rightsquigarrow \{Q'(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma$ and $P <: P'(\bar{x})[\sigma]$.*

$$\begin{array}{c}
\frac{w \models_z F \quad w \models_z F_1}{w \models_z F \wedge F_1} \quad \frac{w \models_{-z} A^g}{w \models_z \neg A^g} \\
\frac{A^g \in w}{w \models_+ A^g} \quad \frac{A^g \notin w}{w \models_- A^g}
\end{array}$$

Figure 7: Declarative interpretation of PDDL formulae. We define $\neg+$ = \neg and $\neg-$ = $+$.

LEMMA 7 (PROPERTY OF STRUCTURAL RULES (RIGHT)). *If there is a derivation $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid \alpha$ by the rules of Figure 6, we have $\{P'(\bar{x})\} \rightsquigarrow \{Q'(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma$ and $Q <: Q'(\bar{x})[\sigma]$.*

Given a planning context Γ , we say that a plan f is *well-typed* (for $\{P\} \rightsquigarrow \{Q\}$), if there is a derivation of $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ by the rules of Figure 6.

3 Soundness of the PCP Logic

We now show that the PCP logic we introduced in previous sections is sound relative to the possible world semantics of PDDL [14].

3.1 Possible World Semantics for PDDL Languages

Coming back to our running example of a PDDL domain, given in Figure 2, we notice that it is defined in a subset of first-order logic, while the actual problem description (Figure 1) contains only ground terms. This motivates us to formally define *PDDL formulae* as follows:

DEFINITION 3 (PDDL FORMULAE).

$$\begin{array}{ll}
\text{Ground Atoms} & \text{GAtom} \ni A^g \quad := R(c_1, \dots, c_n) \\
\text{PDDL Formulae} & \text{Form} \ni F, F_1 \dots F_n \quad := A^g \mid \neg A^g \mid F \wedge F_1
\end{array}$$

Possible world semantics for PDDL [14] is defined in Figure 7. A possible world, or just a *world* is a set of ground atomic formulae. We use letter w to denote a single possible world. Given a world w , a PDDL formula F is *satisfied* by w if $w \models_+ F$ can be derived by

the rules of Figure 7. It should be noted that negation can only be applied to atomic formulae.

It is useful to establish a correspondence between states and formulae. Following [45, 46], we achieve this by introducing a “normalisation” function from PDDL formulae to states.

DEFINITION 4 (NORMALISATION OF PDDL FORMULAE TO STATES[45]). *The function \downarrow_z normalises a PDDL formula to a state:*

$$\begin{aligned} (F \wedge F_1) \downarrow_z S &= F_1 \downarrow_z F \downarrow_z S \\ \neg A^g \downarrow_z S &= A^g \downarrow_{-z} S \\ A^g \downarrow_z S &= A^g \mapsto z * S \end{aligned}$$

We write $F \downarrow_z$ to mean $(F \downarrow_z \text{emp})$.

EXAMPLE 5 (NORMALISATION OF A FORMULA TO A STATE).

$(\text{handEmpty} \wedge \neg \text{onTable a}) \downarrow_+ = \text{handEmpty} \mapsto + * \text{onTable a} \mapsto -$

Normalisation is sound relative to the possible world semantics. A world w_S is a *well-formed world* for a given state S , if the world w_S contains all A^g 's such that $(A^g \mapsto +) \in S$ and contains no A^g 's such that $(A^g \mapsto -) \in S^1$. Generally w_S is not uniquely defined, and we use the notation $\langle w_S \rangle$ to refer to the (necessarily finite) set of all w_S .

EXAMPLE 6 (WELL-FORMED WORLDS).

If $S = (\text{handEmpty} \wedge \neg \text{onTable a}) \downarrow_+$, then w_S may be given by e.g. $w_1 = \{\text{handEmpty}\}$, or $w_2 = \{\text{handEmpty}, \text{onTable b}\}$, or any other world containing handEmpty but not onTable a . The given formula will be satisfied by any such w_S .

Well-formed worlds have the following property :

LEMMA 8 (SUBTYPING AND WELL-FORMED WORLDS). *If we have states P and Q , $Q <: P$ and $w \in \langle w_Q \rangle$ then $w \in \langle w_P \rangle$.*

Finally, we prove that normalisation is sound and complete:

THEOREM 9 (SOUNDNESS AND COMPLETENESS OF NORMALISATION [45, 46]). *Given a formula F and a world w , it holds that $w \models_z F$ iff $w \in \langle w_{F \downarrow_z} \rangle$.*

PROOF. (\Rightarrow) is proven by induction on the derivation of $w \models_z F$. (\Leftarrow) follows by induction on the shape of F , cf. the attached Agda file [20] for the fully formalised proof. \square

3.2 Soundness Theorem

We want to show that if we derive $\Gamma \vdash \{F \downarrow_z\} \rightsquigarrow \{F_1 \downarrow_z\} \mid f$ using the rules given in Figure 6 then we are guaranteed that the evaluation of the plan f on a world that satisfies F produces a new world satisfying F_1 .

To evaluate a plan we will define an *evaluation function* $\llbracket \cdot \rrbracket$ that will interpret actions on worlds. Recall that every state S maps to a world w_S . Let us use notation δ for an arbitrary mapping (an *action handler*) that maps each action $\phi; \{P\} \rightsquigarrow \{Q\} \mid \alpha$ to insertions and deletions on the world w_S according to α 's action on S . We then define the evaluation function $(\llbracket f \rrbracket)^\delta w$ that *evaluates* a plan f in a world w using an action handler δ :

¹By abuse of notation that will not cause confusion, we will use the symbol \in to denote State membership, as well as set membership.

DEFINITION 5 (EVALUATION FUNCTION).

$$\begin{aligned} \llbracket \text{shrink} \rrbracket^\delta w &= w \\ \llbracket a \rrbracket^\delta w &= \delta a w \\ \llbracket f; f_1 \rrbracket^\delta w &= \llbracket f_1 \rrbracket^\delta (\llbracket f \rrbracket^\delta w) \end{aligned}$$

The evaluation function has three cases. The shrink case just returns the world itself, as there is no computational meaning for a shrink action in evaluation. For a single action, evaluation applies the action handler to the world. For a complex plan, evaluation recurses to sub-plans.

The following property of action handlers will be used in the soundness proof:

LEMMA 10 (ACTION HANDLER STRENGTHENING). *If $(\delta \alpha w) \in \langle w_Q \rangle$ and $(\delta \alpha w) \in \langle w_{A \mapsto z} \rangle$ then $(\delta \alpha w) \in \langle w_{A \mapsto z * Q} \rangle$.*

We now proceed to define the notion of a well-formed handler, that will be used to prove soundness of the PCP logic.

DEFINITION 6 (WELL-FORMED HANDLER). *We say that an action handler δ is well-formed if, given:*

- a context Γ with $\phi(\bar{x}); \{P'(\bar{x})\} \rightsquigarrow \{Q(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma$,
- a state P , such that $P(x) <: P'(\bar{x})[\sigma]$ for some ground substitution σ and $\phi(\bar{x})[\sigma]$ normalises to \top ,
- a world $w \in \langle wp \rangle$,

δ satisfies the following property: $(\delta(\alpha(\bar{x})[\sigma]) w) \in \langle w_{P \sqcup Q(\bar{x})[\sigma]} \rangle$.

The next two theorems show that executing a well-typed plan f by the evaluation function $\llbracket f \rrbracket^\delta w$ is *sound*, for any well-formed handler δ .

THEOREM 11 (SOUNDNESS OF EVALUATION FOR NORMALIZED FORMULAE). *Suppose $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$. Then for any $w \in \langle wp \rangle$, and any well-formed handler δ , it follows that $\llbracket f \rrbracket^\delta w \in \langle w_Q \rangle$.*

PROOF. The proof proceeds by structural induction on the typing derivation $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$. In each of the below cases, we take $P, w \in \langle wp \rangle$, and assume $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ was proven by application of a given rule in Figure 6, and in each case we will aim to show that $\llbracket f \rrbracket^\delta w \in \langle w_Q \rangle$.

BASE CASE 1 (ApplyAction). *Suppose we have a proof for $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ by means of the rule ApplyAction. The rules premise requires that some $\phi(\bar{x}); \{P'(\bar{x})\} \rightsquigarrow \{Q'(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma$, and moreover there exists σ s.t. $P'(\bar{x})[\sigma] \equiv P$, $Q'(\bar{x})[\sigma] \equiv Q$, $\alpha(\bar{x})[\sigma] \equiv f$ and $\phi(\bar{x})[\sigma]$ normalises to \top .*

Because δ is well-formed and $w \in \langle wp \rangle$, we have: $(\delta f w) \in \langle w_{P \sqcup Q} \rangle$. We note that $P <: P'(\bar{x})[\sigma]$ because $P \equiv P'(\bar{x})[\sigma]$ by the conditions of the rule, and $P <: P$ by reflexivity of subtyping relation.

It remains to show that $(\delta f w) \in \langle w_{P \sqcup Q} \rangle$ implies that $(\delta f w) \in \langle w_Q \rangle$. We know that $(P \sqcup Q) <: Q$ from Lemma 4 and can therefore deduce $(\delta a w) \in \langle w_Q \rangle$ by applying Lemma 8.

INDUCTIVE CASE 1 (Weakening). *Taking $P, w \in \langle wp \rangle$ as before, we assume $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ was proven by applying Weakening. By inductive hypothesis we know that there is a proof of $\Gamma \vdash \{P'\} \rightsquigarrow \{Q\} \mid f$, such that $P <: P'$ and $\llbracket f \rrbracket^\delta w' \in \langle w_Q \rangle$ if $w' \in \langle wp' \rangle$ for some w' . By Lemma 8, we know that $w \in \langle wp \rangle$ implies $w \in \langle wp' \rangle$. And so we have $\llbracket f \rrbracket^\delta w \in \langle w_Q \rangle$ as required.*

INDUCTIVE CASE 2 (Shrink). We now assume that $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ is obtained by application of *Shrink*, i.e. $f \equiv (f_1; \text{shrink})$ for some f_1 . By inductive hypothesis we know that there is a proof of $\Gamma \vdash \{P\} \rightsquigarrow \{Q'\} \mid f_1$ such that $Q' <: Q$, and $\llbracket f_1 \rrbracket^\delta w \in \langle w_{Q'} \rangle$ if $w \in \langle w_P \rangle$. Because we already have $w \in \langle w_P \rangle$ among our assumptions, we get $\llbracket f_1 \rrbracket^\delta w \in \langle w_{Q'} \rangle$. We apply Lemma 8 to get $\llbracket f_1 \rrbracket^\delta w \in \langle w_Q \rangle$. It remains to show that $\llbracket f_1; \text{shrink} \rrbracket^\delta w \in \langle w_Q \rangle$. By definition of the evaluation function, $\llbracket f_1; \text{shrink} \rrbracket^\delta w = \llbracket f_1 \rrbracket^\delta (\llbracket \text{shrink} \rrbracket^\delta w) = \llbracket f_1 \rrbracket^\delta w$, as required.

INDUCTIVE CASE 3 (Composition). We now assume that $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ by application of *Composition*. By inductive hypothesis we know that, for some f_1 and f_2 such that $f \equiv f_1; f_2$, and for some Q' and Q'' such that $Q' <: Q''$,

- there is a proof of $\Gamma \vdash \{P\} \rightsquigarrow \{Q'\} \mid f_1$ and $\llbracket f_1 \rrbracket^\delta w \in \langle w_{Q'} \rangle$ if $w \in \langle w_P \rangle$;
- there is a proof of $\Gamma \vdash \{Q''\} \rightsquigarrow \{Q\} \mid f_2$ and $\llbracket f_2 \rrbracket^\delta w' \in \langle w_Q \rangle$ if $w' \in \langle w_{Q''} \rangle$;

Because we already have $w \in \langle w_P \rangle$ among our assumptions, we get $\llbracket f_1 \rrbracket^\delta w \in \langle w_{Q'} \rangle$. Next, we need to apply Lemma 8 and the fact that $Q <: Q''$ to get $\llbracket f_1 \rrbracket^\delta w \in \langle w_{Q''} \rangle$. Thus we found a suitable $w' \equiv \llbracket f_1 \rrbracket^\delta w$. But then we get $\llbracket f_2 \rrbracket^\delta (\llbracket f_1 \rrbracket^\delta w) \in \langle w_Q \rangle$. Finally, by definition of the evaluation function, we know that $\llbracket f_1; f_2 \rrbracket^\delta w = \llbracket f_2 \rrbracket^\delta (\llbracket f_1 \rrbracket^\delta w)$. And so we get $\llbracket f_1; f_2 \rrbracket^\delta w \in \langle w_Q \rangle$.

INDUCTIVE CASE 4 (Frame). We now assume that $\Gamma \vdash \{P\} \rightsquigarrow \{Q\} \mid f$ by application of the *Frame* rule, that is, $f \equiv \alpha, P \equiv (P' * A \mapsto z)$, $Q \equiv (Q' * A \mapsto z)$ (for some α, P', Q', A and z), moreover $w \in \langle w_{A \mapsto z * P'} \rangle$, $A \notin P'$, $A \notin Q'$. By the inductive hypothesis, we know that there is a proof of $\Gamma \vdash \{P'\} \rightsquigarrow \{Q'\} \mid \alpha$ and $\llbracket \alpha \rrbracket^\delta w' \in \langle w_{Q'} \rangle$ if $w' \in \langle w_{P'} \rangle$, for any w' .

By Lemma 3 and the fact that $P' <: P'$, we get $A \mapsto z * P' <: P'$. We then use Lemma 8, and our assumption $w \in \langle w_P \rangle$ to assert that $w \in \langle w_{P'} \rangle$, and therefore we get $\llbracket \alpha \rrbracket^\delta w \in \langle w_{Q'} \rangle$. It remains to show that $\llbracket \alpha \rrbracket^\delta w \in \langle w_Q \rangle$.

By the definition of evaluation function, $\llbracket \alpha \rrbracket^\delta w = \delta \alpha w$. Lemma 10 lets us combine two results: 1. $(\delta \alpha w) \in \langle w_{Q'} \rangle$ and 2. $(\delta \alpha w) \in \langle w_{A \mapsto z} \rangle$ to produce the goal $(\delta \alpha w) \in \langle w_{A \mapsto z * Q'} \rangle$ which gives us $(\delta \alpha w) \in \langle w_Q \rangle$ and therefore $\llbracket \alpha \rrbracket^\delta w \in \langle w_Q \rangle$, as required.

It only remains to show that $(\delta \alpha w) \in \langle w_{A \mapsto z} \rangle$. To prove this, we will use the fact that δ is a well-formed handler, and consider $(\delta \alpha w)$. Recall that

- $w \in \langle w_{A \mapsto z * P'} \rangle$, and,
- by inductive hypothesis, there is a derivation for $\Gamma \vdash \{P'\} \rightsquigarrow \{Q'\} \mid \alpha$. Therefore, there is $\phi(\bar{x}); \{P''(\bar{x})\} \rightsquigarrow \{Q''(\bar{x})\} \mid \alpha(\bar{x}) \in \Gamma$ by Lemma 6.
- Also by Lemma 6, we have $P' <: P''(\bar{x})[\sigma]$, for some σ .
- We know that $\alpha(\bar{x})[\sigma]$ must normalise to \top , or there would be no derivation for $\Gamma \vdash \{P'\} \rightsquigarrow \{Q'\} \mid \alpha$.

Given these four conditions, a well-formed handler must satisfy the property: $(\delta \alpha w) \in \langle w_{(A \mapsto z * P') \sqcup Q''(\bar{x})[\sigma]} \rangle$. We can apply Lemma 8 and show that $(\delta \alpha w) \in \langle w_{A \mapsto z} \rangle$, if we can show that $(A \mapsto z * P') \sqcup Q''(\bar{x})[\sigma] <: (A \mapsto z)$. Using Lemma 5 we can establish that $A \mapsto z \in (A \mapsto z * P') \sqcup Q''(\bar{x})[\sigma]$ if $A \notin Q''(\bar{x})[\sigma]$. To show $A \notin Q''(\bar{x})[\sigma]$, we use Lemma 2, Lemma 7 (which gives us $Q' <:$

$Q''(\bar{x})[\sigma]$) and the assumption that $A \notin Q'$. From $A \mapsto z \in (A \mapsto z * P') \sqcup Q''(\bar{x})[\sigma]$ we obtain $(A \mapsto z * P') \sqcup Q''(\bar{x})[\sigma] <: (A \mapsto z)$ by using the subtyping derivation rules. \square

THEOREM 12 (SOUNDNESS OF EVALUATION). Suppose $\Gamma \vdash F_1 \downarrow_+ \rightsquigarrow F_2 \downarrow_+ \mid f$ then for any w such that $w \models_+ F_1$, and any well-formed δ it follows $\llbracket f \rrbracket^\delta w \models_+ F_2$.

PROOF. By assumption $w \models_+ F_1$ and by Theorem 9, we have $w \in \langle w_{F_1 \downarrow_+} \rangle$. Then from Theorem 11, we have $\llbracket f \rrbracket^\delta w \in \langle w_{F_2 \downarrow_+} \rangle$. Thus by Theorem 9, we obtain $\llbracket f \rrbracket^\delta w \models_+ F_2$. \square

Thus if f is well-typed, we are guaranteed that the execution of f in world w is correct.

4 Implementation and Evaluation

As mentioned already, the PCP logic and all lemmas and theorems presented in this paper are formalised in Agda, see [20]. This gives us assurance of the correctness of the presented approach. This Agda module also serves as a *standard library* for verifying PDDL plans. Recall that in Section 1.1, for example, our task was to verify an exact plan f_{ab} , i.e. to derive $\Gamma_{BW} \vdash P_{ab} \rightsquigarrow Q_{ab} \mid f_{ab}$. To do this, we need to create an additional file that defines $\Gamma_{BW}, P_{ab}, Q_{ab}$ and f_{ab} in Agda syntax. Then we need to construct a proof in Agda that this plan is indeed valid. That is, judgements like $\Gamma_{BW} \vdash P_{ab} \rightsquigarrow Q_{ab} \mid f_{ab}$ are not automatically type-checked by Agda, but require manual proofs (using the rules of the PCP logic, cf. Figure 6).

To mitigate this, we automate the following two tasks:

- (1) conversion from PDDL syntax to Agda.

For example the `pickup_from_table` x action as given in Figure 2 is translated to the following snippet of Agda code:

```
 $\Gamma_1$  (pickup_from_table x) = [],
(+, handempty) * (+, ontable x) * (+, clear x) * [],
(+, clear x) * (-, handempty) *
(-, ontable x) * (+, holding x) * []
```

- (2) PCP logic proof generation for Agda, given a PDDL plan.

Figure 8 shows an example of the Agda proof for $\Gamma_{BW} \vdash P_{ab} \rightsquigarrow Q_{ab} \mid f_{ab}$ generated automatically given the domain and planning problem specifications of Figures 1 and 2.

For the first task, we translate a given planning problem and domain to a single Agda file. Conversion of *objects* is one-to-one, i.e. the list of *objects* is given as a list of constructors for the datatype C that stores constants. To convert states, we change the list syntax from Lisp style to Agda style and add the relevant polarity. *Predicates* and *actions* are translated to Agda by representing them as functions from constants to the relevant type. For example the predicate $(\text{on } ?x \ ?y)$ is translated to $\text{on} : C \rightarrow C \rightarrow R$. Action descriptions are described by a parametrised *precondition* and *effect* list in PDDL as shown in Figure 2. The PDDL *precondition* list contains *constraints* and formulas which are separate in our context so we separate them when translating to Agda. Preconditions, constraints and effects are then mapped one-to-one into a context description.

```

P = (+ , (ontable a)) * (+ , (ontable b)) * (+ , (clear a)) * (+ , (clear b)) * (+ , (handempty)) * []

Q : State
Q = (+ , (on a b)) * (+ , (ontable b)) * []

P' : State
P' = (+ , ontable b) * (+ , clear b) * (+ , handempty) * (+ , ontable a) * (+ , clear a) * []

plan : f
plan = (join (join (act (pickup_from_table a)) (act (putdown_on_stack a b))) shrink)

Derivation : Γ1 , P  $\rightsquigarrow$  Q | plan
Derivation = weakening P (from-yes (decSub P' P)) tt (shrink Q tt
  (from-yes (decSub Q ((- , ontable a) * (+ , ontable b) * (+ , clear a) * (- , holding a) *
    (- , clear b) * (+ , on a b) * (+ , handempty) * [])))
  (composition (from-yes (decSub
    ((- , ontable a) * (+ , ontable b) * (+ , clear a) * (+ , holding a) * (+ , clear b) * []))
    ((+ , ontable b) * (+ , clear b) * (+ , clear a) * (- , handempty) * (- , ontable a) * (+ , holding a) * []))))
  ((frame + (ontable b) ( $\lambda z \rightarrow z$ ) ( $\lambda z \rightarrow z$ )
    (frame + (clear b) ( $\lambda z \rightarrow z$ ) ( $\lambda z \rightarrow z$ ) (applyAction tt tt tt))))
  ((frame - (ontable a) ( $\lambda z \rightarrow z$ ) ( $\lambda z \rightarrow z$ )
    (frame + (ontable b) ( $\lambda z \rightarrow z$ ) ( $\lambda z \rightarrow z$ )
      (frame + (clear a) ( $\lambda z \rightarrow z$ ) ( $\lambda z \rightarrow z$ ) (applyAction tt tt tt))))))

```

Figure 8: Agda typing derivation for BlocksWorld problem and domain given in Figures 1 and 2.

For the second task, we implemented a solver for generating Agda type derivations, given a plan. A high level overview of the solver algorithm is shown in Figure 9.

We thus obtain a parser and a proof generator (implemented in Lisp²) that can process plans given in PDDL. However, we delegate the proof-checking (as type-checking) to Agda. This latter step ultimately ensures fully formal plan verification. We call the resulting tool *plan verifier*. The actual implementation [20] contains further instructions and examples.

4.1 Evaluation of the Library Performance

Table 1 shows the results of evaluating the plan verifier over a few benchmark PDDL domains: BlocksWorld, Logistics, Satellite and Mprime[2]. All domains use the STRIPS requirement with Mprime also requiring equality and negative preconditions. All of these examples are generated automatically by supplying a plan and the PDDL domain description to the plan verifier.

This evaluation shows that our system scales from BlocksWorld to more complicated domains, even with increasing plan length. Firstly, it helps to off-load time-consuming plan search to STRIPS. Transforming plans into Agda proofs does not take long (cf. middle column of Table 1). Type-checking time may look worrying, however there is plenty of room for improving it. The given type-checking times reflect the fact that our Lisp script generates excessively long Agda proofs. This happens because we frame all formulas in when we generate the proofs (see Figure 9). For example, for

²Both PDDL and Emacs are written in Lisp, which determined our choice for using Lisp here.

PDDL main	Do-	Plan Length (number of actions)	Proof Generation Time (seconds)	Typechecking Time (seconds)
Blocksworld		10	0.03	10.33
Logistics		24	0.07	28.86
Satellite		9	0.03	15.66
Mprime		11	0.09	42.02

Table 1: Evaluation of the plan verifier. All tests were performed on an Intel Core i5-4670K processor with 8GB of RAM.

the Logistics domain, we have a PDDL plan of length 24. For it, we have PCP/Agda proof with nearly 900 rule applications. Yet, looking closer, we can see that they are mostly frame rules (838 Frame rules, 23 Compositions, 24 ApplyAction, 1 Weakening, 1 shrink).

Similarly, for Mprime example, we have a PDDL plan of length 11, but nearly 900 rule applications in PCP/Agda. Once again, most of them are applications of the frame rule: (865 frame rules, 10 compositions, 11 ApplyAction, 1 weak and 1 shrink).

Ignoring the redundant applications of the frame rule, we see linear dependency of the PCP/Agda proof size relative to STRIPS plan size. Thus, we believe that the type checking time shown in Table 1 does not point to limitations of a type-based approach, but merely reflects the inefficiency of the Lisp script that generates Agda proofs. We will address these shortcomings in future work.

4.2 Leveraging the Power of Dependent Types

Agda is of course also a dependently-typed language. And, as we mentioned in the introduction, the benefit of this approach is the ability to use plans as functions (using the action handlers). One benefit of this would be easy extensions to practical scenarios in which dependent types impose further restrictions and checks on action handlers. Action handlers currently have a type $\text{Action} \rightarrow \text{World} \rightarrow \text{World}$. Within a dependently-typed setting, it is easy to extend this type with say an energy constraint that limits the number of actions that can be taken. Assume a scenario when a robot is given a certain amount of energy, or “fuel”, and each action execution consumes one unit of energy; the robot may not consume more energy than given. It only takes a few lines of code to add this information to our current implementation:

EXAMPLE 7 (ACTION HANDLER ENERGY CONSUMPTION).

```
data Energy : Nat -> Set where
  en : (n : Nat) -> Energy n

EnergyValue : ∀ {n} -> Energy n -> Nat
EnergyValue {n} x = n

actionHandler : Set
actionHandler = ∀ {n} -> Action -> World × Energy (suc n)
               -> World × Energy n
```

Now our implementation incorporates constraints on energy consumption, and the action execution will be bound by the amount of the given energy. This is a really powerful way to use dependent types as it improves readability, and also provides endless possibilities for incorporating various computational constraints in the plan execution.

4.3 Extraction of Plans to Executable Code

We can go one step further, and use Agda’s code extraction library and compile our verified plans into executable Haskell programs or executable byte code. The latter may be deployed directly in robots. The process is fully automated by existing Agda libraries, and subsequent execution of the byte code takes just seconds. For example, we have compiled the BlocksWorld, Logistics and Satellite examples into byte code where all examples run in just 0.02 seconds. We refer the reader to [20] for further details.

4.4 Lessons Learnt: Effects and States

As it often happens with verification projects, this work helps to uncover some previously unknown or unnoticed properties of PDDL. We will give two examples here.

As seen in Figure 2, the syntax of PDDL defines actions by “preconditions” and “effects”. The PCP logic formalises both as states. Yet, there is a subtle difference between an effect and a state. Recall that an effect is executed by deleting all false formula maps from a state and adding all true formula maps. To convert an effect to a state, we must keep the list of all unaffected formula maps intact.

Given a PDDL domain D , a PDDL problem description P_D , and a plan f_1 , the Agda proof script for f_1 in PCP logic is generated as follows:

- (1) Parse D, P_D into Lisp syntax:
 - (a) Store the objects in D , initial and goal world in P_D as variables in Lisp. The initial world id stored in a variable w representing the current world.
 - (b) Convert actions from D into parametrised Lisp functions that generate preconditions and postconditions.
- (2) For all actions in the plan f_1 :
 - (a) Store w in a backup variable so that the Agda subtyping relations can be generated.
 - (b) Use the relevant Lisp function (as defined in (1.b)) to generate the preconditions and postconditions of the current action.
 - (c) Generate the frame axioms by comparing the preconditions of the action to w where all formula maps in w that are not in the preconditions are framed in.
 - (d) Use Lisp functions to apply the action to the world and store the result in the world variable.
- (3) Use all stored results to generate and write Agda proof to file:
 - (a) Start derivation with the Weakening rule to allow for the reordering of the initial state.
 - (b) Use the Shrink rule over the rest of the derivation to shrink the result to the goal state.
 - (c) The rest of the derivation proceeds by composing all actions in the plan f_1 with the relevant subtyping relations and frame axioms.
- (4) Typecheck the generated Agda file to confirm the validity of the proof.

Figure 9: Overview of the code that automatically generates PCP proofs in Agda given PDDL Domain and plan. The code is given in [20].

Also, as we have shown, the states come with the notion of ordering, but effects do not. These simple observations have surprisingly powerful consequences.

Example 1. Ordering and consistency. Take the action *move* from the Logistic domain:

$$[]; \left\{ \begin{array}{l} \text{isVehicle } v \mapsto + \\ * \text{isLocation } loc1 \mapsto + \\ * \text{isLocation } loc2 \mapsto + \\ * \text{isAt } v \text{ } loc1 \mapsto + \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{l} \text{isVehicle } v \mapsto + \\ * \text{isLocation } loc1 \mapsto + \\ * \text{isLocation } loc2 \mapsto + \\ * \text{isAt } v \text{ } loc1 \mapsto - \\ * \text{isAt } v \text{ } loc2 \mapsto + \end{array} \right\}$$

Imagine we instantiated the move action with (car museum museum). In the PCP logic, this instantiation will produce an inconsistent state:

$$\left\{ \begin{array}{l} \text{isVehicle car} \mapsto + \\ * \text{isLocation museum} \mapsto + \\ * \text{isLocation museum} \mapsto + \\ * \text{isAt car museum} \mapsto + \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{l} \text{isVehicle car} \mapsto + \\ * \text{isLocation museum} \mapsto + \\ * \text{isLocation museum} \mapsto + \\ * \text{isAt car museum} \mapsto - \\ * \text{isAt car museum} \mapsto + \end{array} \right\}$$

In PDDL the effect will simply be executed. The result of this action will depend on the order in which the effect formulas are executed. And, since PDDL specifications [14] do not specify any particular ordering on effect formulas, planners have to make this decision themselves. So, some planners come to the conclusion that the car is at the museum, and some – that it is not.

In the PCP logic, this plan will simply not be type-checked and the user will receive a due type checking error.

Example 2. Loss in Translation In our early experiments, we encountered a problem that many good plans are not type-checked when they are translated verbatim to the PCP logic. The reason for this is the loss of information between the “precondition” and the “effect” in the PDDL formulation. We use the following example to explain the problem.

Consider the `pickup_from_stack` action from the BlockWorld domain definition:

$$[]; \left\{ \begin{array}{l} \text{on } x_1 \ x_2 \mapsto + \\ * \text{clear } x_1 \mapsto + \\ * \text{handEmpty} \mapsto + \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{l} \text{on } x_1 \ x_2 \mapsto - \\ * \text{handEmpty} \mapsto - \\ * \text{holding } x_1 \mapsto + \\ * \text{clear } x_2 \mapsto + \end{array} \right\}$$

Notice that `clear x_1` $\mapsto +$ is not mentioned in the effect list, because this fact is unaffected by the action. But if we treat this as a state, rather than effect, the information about `clear x_1` $\mapsto +$ will simply be lost. In the PCP logic, the frame rule can not be used to recover this information, as this formula already occurs in the precondition. As a result, some PDDL plans will fail to type check in the PCP logic. To fix this problem, we add all such formula maps explicitly to the postconditions:

$$[]; \left\{ \begin{array}{l} \text{on } x_1 \ x_2 \mapsto + \\ * \text{clear } x_1 \mapsto + \\ * \text{handEmpty} \mapsto + \end{array} \right\} \rightsquigarrow \left\{ \begin{array}{l} \text{on } x_1 \ x_2 \mapsto - \\ * \text{handEmpty} \mapsto - \\ * \text{holding } x_1 \mapsto + \\ * \text{clear } x_2 \mapsto + \\ * \text{clear } x_1 \mapsto + \end{array} \right\}$$

The plan verifier we implement does this transformation automatically.

5 Conclusions, Related and Future Work

We have presented the PCP logic, a novel resource logic for verification of AI plans, and proven its soundness relative to the possible world semantics of PDDL. We have shown the benefits that resource semantics and the Curry-Howard correspondence bring to this framework. In particular, the former makes it easier to formalise state consistency and other constraints within the logic, and the latter enables direct deployment of verified plans as functions. We also presented an Agda library in which the soundness result is proven, and which simultaneously serves as a generic module for verifying plans produced by AI planning. To further strengthen the practical significance of these results, we implemented scripts for automated parsing of PDDL plans, and for automated generation of proofs of their soundness in the PCP logic. The ultimate

proof- (and type-) checking of these is delegated to the Agda library. We evaluated this implementation on several famous PDDL benchmarks.

Our Earlier Work on Proof Carrying Plans. Compared to our earlier attempts to define a “proof-carrying plans” approach [45], this new attempt is stronger both theoretically and practically. The new PCP logic takes inspiration from resource logics as a consequence provides a more natural way to perform local reasoning. This, in its turn, helps to verify not just the plans, but also consistency of domains and states. In previous work the consistency assumption was needed to be stated as an axiom in order to prove soundness of the system, and was not incorporated into checking of individual plans. The PCP Logic embeds consistency directly into the system through its rules. This has two advantages. The first is that it is impossible to derive proofs that contain inconsistent states and the second is that type errors for inconsistency will show exactly where and why there is an inconsistent state. The PCP logic also enables extensions to first-order logic, introduction of richer verification constraints, and opens the possibilities for extensions to concurrent logics. Though the latter extension is left as future work.

From the practical point of view, the earlier work contained no automation presented here. Also, it did not include reasoning with constraints, or any experiments with using the dependent types during the plan execution.

Origins of the Frame Rule. The “frame problem” that inspired the frame rule of Separation logic actually has origins in AI [9, 18]. Initially, the problem referred to the difficulty in local reasoning about problems in a complex world. In AI planning specifically, this problem consisted of keeping track of the consequences of applying an action on a world. Intuitively a person would understand picking up some block a that is on the table would have no effect on some other block b that is on the table. The frame problem deals with the way to represent this intuition formally.

One way to deal with the frame problem is to declare “frame axioms” for every action explicitly. This is an inefficient way to deal with this problem as defining these frame axioms becomes infeasible the larger the system gets [9]. Since most actions in AI planning only make small local changes to the world, a more general representation would be more suitable. STRIPS deals with this problem by introducing an assumption that every formula in a world that is not mentioned in the effect list of an action remains the same after execution of the action. This is known as the “STRIPS assumption” and it is an assumption that PDDL also uses.

The logic of Bunched Implications [24, 35] and Separation Logic [36] took inspiration from this older notion of the frame problem, and introduced more abstract formalism, which is now known as a “frame rule”, into the resource logics [41]. This family of logics has brought many theoretical and practical advances to modelling of complex systems, and is behind many *lightweight verification* projects [6].

In this paper, we have shown how the original frame problem from AI maps back to the more abstract ideas of resource logics. We see this as one of the paper’s contributions.

Curry-Howard Approaches to Separation Logic and Other Resource Logics. The PCP Logic introduces a Curry-Howard approach to AI planning inspired by resource logic. This is in part

inspired by existing applications of the Curry-Howard approach in the field. Both Hoare logic and Separation logic have been given a Curry-Howard interpretation: [31, 33]. Several papers explore the computational and practical benefits of it. For example, Polikarpova and Sergey [40] took a Curry-Howard approach to Separation logic to improve program synthesis seen as a proof search problem. In a similar way to our specifications, they define a synthesis goal $\Gamma \vdash P \rightsquigarrow Q$, which is solved by a program c if the assertion $\Gamma \vdash P \rightsquigarrow Q|c$ can be derived in their system.

In this paper, we also make an attempt to make a case for computational and practical uses of Curry-Howard interpretation of the newly introduced PCP logic.

AI Planning and Linear Logic. There is a long history of modelling AI planning in Linear logic, that dates back to the 90s [25], and was investigated in detail in the 2000s, see e.g. [8, 47]. In fact, AI planning is used as one of the iconic use-cases of Linear logic [39]. The main idea behind using Linear logic for AI planning is treating action descriptions as linear implications:

$$\alpha : \forall x. P \multimap Q,$$

where P and Q are given by tensor products of atoms: $R_1(t_1) \otimes \dots \otimes R_n(t_n)$. We could incorporate information about polarities inside the predicates, as follows: $R_1(t_1, z_1) \otimes \dots \otimes R_n(t_n, z_n)$. Then, the linear implication and the tensor products model the resource semantics of PDDL rather elegantly.

The computational (Curry-Howard) interpretation of AI plans was not the focus of study in the above mentioned approaches, yet it plays a crucial rôle in the PCP logic, from design all the way to implementation, verification and proof extraction (see Section 4).

AI Planning and (Linear) Logic Programming. The above syntax also resembles linear logic programming Lolli, introduced by Miller et al [22]. Lolli was applied in speech planning in [10].

Our previous work [45] in fact takes inspiration from Curry-Howard interpretation of Prolog[15, 16]. In our previous work and in general, logic programming does not work well with PDDL negation. In PDDL, we have to work with essentially three-valued logic: an atom may be declared to be absent or present in a world. But if neither is declared, we assume a “not known” or “either” situation. Logic programming usually uses the approach of “negation-as-failure” that does not agree with this three-valued semantics. A solution is to introduce polarities as terms, as shown in the example above. This merits further investigation.

Curry-Howard view on Linear Logic. Curry-Howard semantics of Linear logic also attracted attention of logicians first in the 90s [3], and then in the 2000s in connection with research into Linear Logical Frameworks [7, 44].

We conjecture that many results obtained in this paper could be replicated in one of these systems. We plan to investigate this approach in comparison with the PCP logic in the future. Generally speaking, the PCP logic can be seen as a domain specific language for AI planning. It is simpler and less expressive than Linear logic but makes up for it in simplicity and close correspondence to PDDL syntax. Transformations between PDDL domain and problem descriptions to the PCP logic are straightforward since the syntax is so similar. This enables us to automate the generation of Agda proofs from PDDL plans. Notably, we have typing rules

for functions that are given directly by PDDL plans. Thus, we verify outputs of PDDL planners as given. This close correspondence to the plans would be impossible in either of the above Curry-Howard versions of Linear logic, where proof terms tend to be much more complex. Pros and cons of domain specific versus general approaches to verification of AI plans deserves further investigation.

We hope that the DSL nature of the PCP logic will pave the way for its wider adoption as a *practical verification tool* for the AI planning community. This is something that previously proposed Linear logic approaches to AI planning did not achieve.

Modelling looping behaviour and non-termination in AI planning. The design of this Agda prototype has revealed several limitations in state-of-the-art implementations of planning languages: e.g. their reliance on the closed world assumption and formulae grounding and the absence of functions. We see the potential of our method to overcome many of these limitations thanks to our general dependently-typed set-up, in which the use of functions, higher-order features, constraints and effect handling will be much more natural than in the current implementations.

Other Future Work. One limitation of the PCP logic is that it only works with a subset of the domains that can be expressed in PDDL. To incorporate more of the PDDL syntax we want to extend the system to reason about temporal (as well as concurrent) planning. We believe that this extension can be naturally expressed in our system due to related extensions in the resource logics.

From the theoretical point of view, we hope to achieve a deeper understanding of the relation of the new PCP logic to the categorical and coalgebraic semantics of other resource logics [41].

We plan to improve the performance of our system, to speed up type checking, and make Agda proof generation more reliable and practical. The former can be improved through the creation of a frame minimising algorithm. The latter can be facilitated by producing partial Agda proofs when the full proof generation is too hard.

Interactive facilities of our tool also deserve future attention. Generally, Agda allows *holes* to be left in a proof which a user can use to interactively inspect the subgoal of the proof. In the future we plan to update our proof generator to generate incomplete proofs so a user can inspect the proof goals that cannot be solved.

Another possibility is to further explore the dependently-typed aspects of our system as described in Section 4.2. This can include extensions such as higher-order functions and universal formulae.

6 Acknowledgements

We thank Simon Docherty, James McKinna and David Pym, for inspiring discussions, constructive criticisms and pointers to related work.

We thank EPSRC DTA PhD Scheme for funding the first author.

The second author acknowledges support of the UK National Cyber Security Center grant *SecCon-NN: Neural Networks with Security Contracts - towards lightweight, modular security for neural networks* and the UK Research Institute in Verified Trustworthy Software Systems (VETSS)-funded research project *CONVENER: Continuous Verification of Neural Networks*.

References

- [1] Mohammad Abdulaziz, Charles Gretton, and Michael Norrish. 2019. A Verified Compositional Algorithm for AI Planning. In *10th International Conference on Interactive Theorem Proving (ITP 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [2] aibasel. 2020. Downward Benchmarks. <https://github.com/aibasel/downward-benchmarks>
- [3] David W. Albrecht, John N. Crossley, and John S. Jeavons. 1997. New Curry-Howard Terms for Full Linear Logic. *Theor. Comput. Sci.* 185, 2 (1997), 217–235.
- [4] Saddek Bensalem, Klaus Havelund, and Andrea Orlandini. 2014. Verification and validation meet planning and scheduling.
- [5] Josh Berdine, Cristiano Calcagno, and Peter W O’hearn. 2005. Symbolic execution with separation logic. In *Asian Symposium on Programming Languages and Systems*. Springer, 52–68.
- [6] Cristiano Calcagno, Dino Distefano, Jérémy Dubreil, Dominik Gabi, Pieter Hooimeijer, Martino Luca, Peter O’Hearn, Irene Papakonstantinou, Jim Purbrick, and Dulma Rodriguez. 2015. Moving fast with software verification. In *NASA Formal Methods Symposium*. Springer, 3–11.
- [7] Ilario Cervesato and Frank Pfenning. 2002. A Linear Logical Framework. *Inf. Comput.* 179, 1 (2002), 19–75.
- [8] Lukáš Chrpa, Pavel Surynek, and Jiří Vyskočil. 2007. Encoding of planning problems and their optimizations in linear logic. In *Applications of Declarative Programming and Knowledge Management*. Springer, 54–68.
- [9] Daniel C Dennett. 2006. Cognitive wheels: The frame problem of AI. (2006).
- [10] Lucas Dixon, Alan Smail, and Tracy Tsang. 2009. Plans, Actions and Dialogues Using Linear Logic. *Journal of Logic, Language and Information* 18, 2 (2009), 251–289.
- [11] George W Ernst and Allen Newell. 1969. *GPS: A case study in generality and problem solving*. Academic Pr.
- [12] Richard Fikes and Nils J. Nilsson. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2, 3/4 (1971), 189–208.
- [13] Kathleen Fisher, John Launchbury, and Raymond Richards. 2017. The HACMS program: using formal methods to eliminate exploitable bugs. *Phil. Trans. Royal Society* (2017), Issue A 375.
- [14] Maria Fox and Derek Long. 2003. PDDL2.1: An extension to PDDL for expressing temporal planning domains. *Journal of artificial intelligence research* 20 (2003), 61–124.
- [15] Peng Fu and Ekaterina Komendantskaya. 2017. Operational semantics of resolution and productivity in Horn clause logic. *Formal Asp. Comput.* 29, 3 (2017), 453–474.
- [16] Peng Fu, Ekaterina Komendantskaya, Tom Schrijvers, and Andrew Pond. 2016. Proof Relevant Corecursive Resolution. In *Functional and Logic Programming - 13th International Symposium, FLOPS 2016, Kochi, Japan, March 4-6, 2016, Proceedings (Lecture Notes in Computer Science)*, Vol. 9613. Springer, 126–143.
- [17] Cordell Green. 1969. Theorem proving by resolution as a basis for question-answering systems. *Machine intelligence* 4 (1969), 183–205.
- [18] Patrick J Hayes. 1981. The frame problem and related problems in artificial intelligence. In *Readings in Artificial Intelligence*. Elsevier, 223–230.
- [19] James A Hendler, Austin Tate, and Mark Drummond. 1990. AI planning: Systems and techniques. *AI magazine* 11, 2 (1990), 61–61.
- [20] A. Hill. 2020. PCP Logic: Agda formalisation. <https://github.com/PDTypes/PCPLLogic>
- [21] Charles Antony Richard Hoare. 1969. An axiomatic basis for computer programming. *Commun. ACM* 12, 10 (1969), 576–580.
- [22] Joshua S. Hodas and Dale Miller. 1994. Logic Programming in a Fragment of Intuitionistic Linear Logic. *Inf. Comput.* 110, 2 (1994), 327–365.
- [23] Richard Howey, Derek Long, and Maria Fox. 2004. VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence*. IEEE, 294–301.
- [24] Samin S Ishtiaq and Peter W O’Hearn. 2001. BI as an assertion language for mutable data structures. In *Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 14–26.
- [25] Eric Jacopin. 1993. Classical AI planning as theorem proving: The case of a fragment of linear logic. In *AAAI Fall Symposium on Automated Deduction in Nonstandard Logics*. AAAI Press Publications Palo Alto, CA, 62–66.
- [26] Xavier Leroy et al. 2012. The CompCert verified compiler. *Documentation and user’s manual*. INRIA Paris-Rocquencourt 53 (2012).
- [27] Derek Long, Maria Fox, and Richard Howey. 2009. Planning domains and plans: validation, verification and analysis. In *Proc. Workshop on V&V of Planning and Scheduling Systems*.
- [28] John McCarthy and Patrick J Hayes. 1981. Some philosophical problems from the standpoint of artificial intelligence. In *Readings in artificial intelligence*. Elsevier, 431–450.
- [29] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. 1998. PDDL-the planning domain definition language. (1998).
- [30] Aleksandar Nanevski. [n.d.]. Separation Logic and Concurrency.
- [31] Aleks Nanevski. [n.d.]. Separation Logic and Concurrency (OPLSS 2016) Draft of July 22, 2016. ([n. d.]).
- [32] Aleksandar Nanevski, Anindya Banerjee, and Deepak Garg. 2013. Dependent type theory for verification of information flow and access control policies. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 35, 2 (2013), 1–41.
- [33] Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. 2006. Polymorphism and separation in hoare type theory. In *Proceedings of the eleventh ACM SIGPLAN international conference on Functional programming*. 62–73.
- [34] George C. Necula. 2011. Proof-Carrying Code. In *Encyclopedia of Cryptography and Security, 2nd Ed*, Henk C. A. van Tilborg and Sushil Jajodia (Eds.). Springer, 984–986.
- [35] Peter O’Hearn, John Reynolds, and Hongseok Yang. 2001. Local reasoning about programs that alter data structures. In *International Workshop on Computer Science Logic*. Springer, 1–19.
- [36] Peter W O’hearn. 2007. Resources, concurrency, and local reasoning. *Theoretical computer science* 375, 1-3 (2007), 271–307.
- [37] J Scott Penberthy, Daniel S Weld, et al. 1992. UCPOP: A Sound, Complete, Partial Order Planner for ADL. *Kr* 92 (1992), 103–114.
- [38] John Penix, Charles Pecheur, and Klaus Havelund. 1998. Using model checking to validate AI planner domain models. In *Proceedings of the 23rd Annual Software Engineering Workshop, NASA Goddard*.
- [39] Jeff Polakow and Frank Pfenning. 2001. *Ordered linear logic and applications*. Carnegie Mellon University Pittsburgh.
- [40] Nadia Polikarpova and Ilya Sergey. 2019. Structuring the synthesis of heap-manipulating programs. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 72.
- [41] David Pym. 2019. Resource semantics: logic as a modelling technology. *ACM SIGLOG News* 6, 2 (2019), 5–41.
- [42] John C Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*. IEEE, 55–74.
- [43] Albert Rizaldi, Fabian Immler, Bastian Schürmann, and Matthias Althoff. 2018. A formally verified motion planner for autonomous vehicles. In *International Symposium on Automated Technology for Verification and Analysis*. Springer, 75–90.
- [44] Anders Schack-Nielsen and Carsten Schürmann. 2008. Celf - A Logical Framework for Deductive and Concurrent Systems (System Description). In *Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings (Lecture Notes in Computer Science)*, Alessandro Armando, Peter Baumgartner, and Gilles Dowek (Eds.), Vol. 5195. Springer, 320–326.
- [45] Christopher Schwaab, Ekaterina Komendantskaya, Alasdair Hill, František Farka, Ronald PA Petrick, Joe Wells, and Kevin Hammond. 2019. Proof-Carrying Plans. In *International Symposium on Practical Aspects of Declarative Languages*. Springer, 204–220.
- [46] C. Schwaab, A. Hill, F. Farka, and E. Komendantskaya. 2018. Proof-Carrying plans: Agda implementation and examples. <https://github.com/PDTypes>
- [47] Mark Steedman. 2002. Plans, affordances, and combinatorial grammar. *Linguistics and Philosophy* 25, 5-6 (2002), 723–753.
- [48] David E Wilkins. 2014. *Practical planning: extending the classical AI planning paradigm*. Elsevier.