

Dynamic Doxastic Differential Dynamic Logic for Belief-Aware Cyber-Physical Systems^{*}

João G. Martins^{1,2} , André Platzer^{1,3} , and João Leite² 

¹ Computer Science Department, Carnegie Mellon University
{jmartins, aplatzer}@cs.cmu.edu

² NOVA LINCS, Universidade NOVA de Lisboa
jleite@fct.unl.pt

³ Fakultät für Informatik, Technische Universität München

Abstract. Cyber-physical systems (CPS), such as airplanes, operate based on sensor and communication data, i.e. on potentially noisy or erroneous beliefs about the world. Realistic CPS models must therefore incorporate the notion of beliefs if they are to provide safety guarantees in practice as well as in theory. To fundamentally address this challenge, this paper introduces a first-principles framework for reasoning about CPS models where control decisions are explicitly driven by controller beliefs arrived at through observation and reasoning. We extend the differential dynamic logic $d\mathcal{L}$ for CPS dynamics with belief modalities, and a learning operator for belief change. This new dynamic doxastic differential dynamic logic $d^4\mathcal{L}$ does due justice to the challenges of CPS verification by having 1) real arithmetic for describing the world and beliefs about the world; 2) continuous and discrete world change; 3) discrete belief change by means of the learning operator. We develop a sound sequent calculus for $d^4\mathcal{L}$, which enables us to illustrate the applicability of $d^4\mathcal{L}$ by proving the safety of a simplified belief-triggered controller for an airplane.

Keywords: differential dynamic logic · dynamic epistemic logic · sequent calculus · hybrid systems · cyber-physical systems

1 Introduction

Cyber-physical systems (CPS) mix discrete cyber change and continuous physical change. Examples of CPS include self-driving cars, airplane autopilots, and industrial machines. With widespread espousal of automation in transportation, it is imperative that we develop methods capable of verifying the safety of the algorithms driving the CPSs on which human lives will increasingly depend.

However, because CPSs rely on sensors and partial human operation, both of which are imperfect, they face a possible discrepancy between reality, and the perception, understanding and beliefs thereof. Critical system components

^{*} Supported by the Alexander von Humboldt Foundation, NSF grant CNS-1446712, CMU | Portugal grant SFRH/BD/51886/2012, and PTDC/CCI-COM/30952/2017

are engineered to be exceptionally reliable, so safety incidents often originate from just such a discrepancy between what is believed to be true versus what is actually true. This can be highlighted by three (of many) tragedies, some now known to be preventable, e.g., through neutral control inputs [5,1,12]. However, non-critical sensor failures led to erroneous pilot beliefs. These beliefs resulted in the pilots’ inability to perform informed, safe control decisions, leading to 574 fatalities in these three incidents alone.

Verification efforts for practical system designs must therefore augment initial analyses which assume perfect information with an awareness of factors such as sensor errors, actuator disturbances, and, crucially, incomplete or incorrect perceptions of the world. Ideally, such factors ought to become an explicit part of the model so that CPS design and verification engineers can confront this challenge of uncertainty head on at design time, before safety violations occur.

We argue that the notion of beliefs (*doxastics*) about the state of the world, which has been extensively studied, can succinctly capture such phenomena. We develop a first-principles language and verification method for reasoning about *changing beliefs in a changing world*. Using this language, CPS designers may create more realistic controllers whose decisions are explicitly driven by their beliefs. The consequences of such decisions are borne out in the continuous-time and continuous-space evolution of these *belief-aware CPS*.

In this new paradigm, control decisions are grounded *only* in what can be observed and reasoned. By providing the tools to develop such *belief-triggered controllers*, we help bridge the gap between the theoretical safety of *CPS models*, and the practical safety of the *CPS vehicles* that will soon be driving and flying us to our destinations.

2 Technical Approach

Our approach is to integrate a framework for specifying and verifying real-world CPS with a suitable notion of dynamic beliefs. The result should be a single cohesive framework capable of complex reasoning about changing beliefs in a changing world, as required by belief-aware CPSs.

Work on control-theoretic *robust* solutions for CPS models seem promising, since they entail asymptotical steering towards a desired target domain despite perturbations in the system [11]: sensor and actuator noise could be modeled as perturbations rather than beliefs. However, perturbation analysis does not capture the complex causal relationship from observation, to reasoning, to actuation in an explicit way that can lead to e.g. malfunction checklists or pilot best practices. Accurate analyses for safety incidents such as [5,1,12] require the power to 1) model agents with reasoning capabilities, and 2) leverage complex logical arguments about perception versus fact in the pursuit of safety guarantees.

The differential dynamic logic $d\mathcal{L}$ [16,17,19] is a successful tool for designing and verifying belief-*unaware* CPS, i.e. a “changing world” in a real-valued domain. Dynamic epistemic logics (DELs), on the other hand, deal with chang-

ing knowledge (which is tightly connected to beliefs⁴) in a propositional static world that never changes [3,4,10,7], again through the lens of modal logic. Some previous work exists at the intersection of these two. However, belief-aware CPS requires unobservable world change under the real numbers, which is in conflict with the public propositional world-change in [6]; and a more comprehensive and less restrictive treatment of belief that goes beyond using the underlying dynamic modalities of world-change to emulate noise as in [14].

Since both \mathbf{dL} and DELs are dynamic modal logics, they are prime candidates for inspiration in the pursuit of a unified dynamic modal logic that can reason about changing beliefs in a changing world. We develop the *dynamic doxastic differential dynamic logic* $\mathbf{d}^4\mathcal{L}$, as an extension of \mathbf{dL} with 1) belief modalities, and 2) a learning operator for describing belief-change, inspired by DELs.

This new framework requires a fundamental conceptual shift in the design of CPS. Let **ctrl** be a program describing control decisions (e.g. a pilot pressing a button), and **plant** be a program for continuous evolution (e.g. an airplane flying). In the current, belief-unaware \mathbf{dL} paradigm, the primary mode of establishing the safety of CPS is by the validity of a formula $pre \rightarrow [(\mathbf{ctrl}; \mathbf{plant})^*] safe$. It states that, starting from precondition pre , every possible execution of the program $(\mathbf{ctrl}; \mathbf{plant})^*$ ends with the safety property $safe$ being true, with the star * operator repeating **ctrl** followed by **plant** any number of times.

Example 1. As a running example, suppose an airplane is controlled by directly setting its vertical velocity to 1 or -1 in thousands of feet per second. The safety goal of the controller is to keep the airplane above ground:

1. $pre \equiv safe \equiv (alt > 0)$, i.e., the airplane is above ground.
2. $\mathbf{ctrl} \equiv (?alt > 1; yv := -1) \cup yv := 1$, in which two things may happen, on either side of \cup . If the airplane is above 1000ft ($?alt > 1$), it may descend by setting vertical velocity yv to -1000 feet per second. Alternatively, it can climb with $yv := 1$, which may always happen since this action has no ? test.
3. $\mathbf{plant} \equiv t := 0; t' = 1, alt' = yv \ \& \ t \leq 1$ describes, using differential equations, that altitude changes with vertical velocity ($alt' = yv$) for a maximum of 1 unit of time using time counter $t' = 1$. The *evolution domain constraint* $t \leq 1$ bounds how much time may pass before the pilot reassesses this choice.

Intuitively, this CPS is safe because the controller can only decide to descend if it is high enough above ground such that descending for 1 second at a velocity of -1000 feet per second, traveling a total of 1000 feet, keeps it above ground. This condition is based on *ontic* (real world, or factual) truth and does not capture the reality that altitude is read from a noisy altimeter, and that *pilot beliefs* trigger actions, not ontics.

In contrast, in *belief-aware* CPS, control decisions are triggered by some belief $B_a(\phi)$, not ontic truth ϕ . This minor syntactic change belies the complexity of the underlying paradigm shift. The CPS model must now explicitly describe how

⁴ Beliefs may be erroneous, knowledge may not.

an agent learns about the world and acquires such beliefs $B_a(\phi)$. In $d^4\mathcal{L}$, this process of observation and reasoning is specified by means of a *learning operator*.

A $d\mathcal{L}$ program α , describing ontic change, does not alter beliefs. In contrast, a learning operator program $L_a(\alpha)$ changes *only* agent a 's beliefs, with the change described by α becoming doxastic rather than ontic. The pattern $\alpha; L_a(\alpha)$ describes *observed* ontic change, which also affects beliefs. This learning operator may be used in a program **obs** to describe the agent's learning processes of observation and reasoning. This leads to the addition the belief-changing **obs** to the safety formula $pre \rightarrow [(\mathbf{obs}; \mathbf{ctrl}; \mathbf{plant})^*] safe$ used for belief-aware CPS.

Example 2. Consider a belief-triggered controller for the airplane of Example 1. The model now incorporates the fact that observation is imperfect, and that the altimeter, while operating properly, has some noise bounded by $\varepsilon > 0$.

1. **obs** $\equiv L_a(?alt - alt_a < \varepsilon)$. The pilot a learns, by observing the altimeter with known error bounds ε , that the *perceived* altitude alt_a can be lower than the *true* altitude alt by at most ε . Thus, the belief $B_a(alt - alt_a < \varepsilon)$ comes to be.
2. **ctrl** $\equiv (?B_a(alt_a - \varepsilon > 1); yv := -1) \cup yv := 1$. Climbing, being safe, remains an always acceptable choice. However, the trigger for descending is that the pilot *believes* that the *perceived* altitude with worst-case noise is still high enough for the airplane to descend for one second, i.e. $B_a(alt_a - \varepsilon > 1)$.

We must add $\varepsilon > 0$ to *pre*, but **plant** does not change since beliefs do not directly affect the behavior of the real world: they do so only through agent actions.

More generally, $d^4\mathcal{L}$ allows for arbitrary combinations of ontic $d\mathcal{L}$ actions and the learning operator, representing any interleaving of physical and doxastic change, the former potentially unobservable, and the latter potentially imperfect, e.g. through noisy sensors.

3 Syntax of $d^4\mathcal{L}$

In this section, we will describe $d^4\mathcal{L}$ terms, formulas and programs. As in $d\mathcal{L}$, real arithmetic is used to accurately model CPSs. Thus, terms are real-valued.

The safety of well-functioning belief-aware CPS is often predicated on beliefs being grounded in reality so that informed decisions can be made, cf. formula $B_a(alt - alt_a < \varepsilon)$ of Example 2, where perceived altitude can underestimate factual altitude by at most ε . This relation between belief and truth is at the core of many safety arguments, and should be describable within the logic. We must therefore be able to refer to both ontic (factual) and doxastic (belief) states in the same context, as in $B_a(alt - alt_a < \varepsilon)$.

3.1 $d^4\mathcal{L}$ Terms and Formulas

State variables describe ontic truth, e.g. alt is the airplane's real altitude. Doxastic variable alt_a is agent a 's perception of alt . Basic arithmetic is also in the

language, e.g. $x - y$. Constants $c \in \mathbb{Q}$ allow for digitally representable numbers in the syntax, e.g. 2.5 but not π , though the semantics can give variables any value in \mathbb{R} . Logical variables X are introduced by quantifiers over \mathbb{R} to e.g. discharge reasoning about continuous time, or to find witnesses for existential modalities.

Let \mathbb{A} be a finite set of agents, Σ be a countable set of logical variables, \mathbb{V} be a countable set of state variables, and $\mathbb{V}_a = \{x_a : x \in \mathbb{V}\}$ the set of doxastic variables for agent $a \in \mathbb{A}$. The following definition distinguishes between terms with and without doxastic variables. The distinction is crucial when assigning to state or doxastic variables, as we will see in Definition 3.

Definition 1. *The doxastic terms θ and non-doxastic terms ζ of $d^4\mathcal{L}$, with $\otimes \in \{+, -, \times, \div\}$, $X \in \Sigma$, $x \in \mathbb{V}$, $x_a \in \mathbb{V}_a$, $a \in \mathbb{A}$, $c \in \mathbb{Q}$, are given by the grammar:*

$$\begin{aligned} \theta & ::= \theta \otimes \theta \mid X \mid c \mid x \mid x_a \\ \zeta & ::= \zeta \otimes \zeta \mid X \mid c \mid x \end{aligned}$$

The formulas of $d^4\mathcal{L}$ are a superset of $d\mathcal{L}$'s [17], which are a superset of those of first-order logic for real arithmetic. Alongside logical connectives, we may write propositions such as $\theta_1 \leq \theta_2$ and logical quantifiers $\forall X \phi$. To this, $d^4\mathcal{L}$ adds the belief modality $B_a(\phi)$, meaning agent a believes ϕ . The dynamic modality formula $[\alpha]\phi$ (after all executions of program α , ϕ is true), and its dual $\langle\alpha\rangle\phi$ (after some execution of α , ϕ is true) capture belief-aware CPS behavior. The language of the programs α will be specified later in Definition 3.

Since $d^4\mathcal{L}$ beliefs are *only* about the state of the world, it is useful to distinguish between formulas ξ which may appear inside belief modalities, and those ϕ which may not. We still allow doxastic terms θ in ϕ , since safety proofs may generate such formulas.

Definition 2. *The formulas ϕ, ξ of $d^4\mathcal{L}$ are given by the grammar:*

$$\begin{aligned} \phi & ::= \phi \vee \phi \mid \neg\phi \mid \theta \leq \theta \mid \forall X \phi(X) \mid [\alpha]\phi \mid B_a(\xi) \\ \xi & ::= \xi \vee \xi \mid \neg\xi \mid \theta \leq \theta \end{aligned}$$

The remaining logical connectives, \wedge , \rightarrow and duals $\langle\alpha\rangle\phi$, $\exists X \phi(X)$, $P_a(\xi)$ are defined as usual, e.g. $\langle\alpha\rangle\phi \equiv \neg[\alpha]\neg\phi$, and $P_a(\xi) \equiv \neg B_a(\neg\xi)$ when a considers ξ possible. We may now generalize the noisy but accurate sensors of Example 2.

Example 3 (Noisy sensors). Sensors often come with known error bounds ε . A pilot reading from the altimeter should thus come to believe the indicated value to be within ε of the real *alt*, as captured by $B_a((alt_a - alt)^2 \leq \varepsilon^2)$, with integer exponentiation being definable from multiplication.

Belief modalities with both state and doxastic variables are *meta-properties* of belief, e.g., how far doxastic truth is from ontic truth. Thus, their truth value indeed changes as either the world or beliefs change. Section 6 will show such formulas are part of the core argument for some belief-aware CPS safety proofs. When formulas such as $B_a((alt_a - alt)^2 \leq \varepsilon^2)$ are not true, it can become impossible for a to make informed decisions. Safety may then instead rely on very conservative actions, e.g. bringing a car to a stop, or flying straight and level.

3.2 Doxastic Hybrid Programs

The hybrid programs (HPs) of $d\mathcal{L}$ [17] are able to describe both discrete and continuous ontic change. They are the starting point for the *doxastic hybrid programs* (DHPs) of $d^4\mathcal{L}$. We introduce a learning operator $L_a(\gamma)$ for doxastic change, where γ encodes an agent observing the world, reading from a sensor, or suspecting some change to have happened. In this paper, the language of the learned program γ is nearly identical to that of hybrid programs, and to the epistemic actions of the epistemic action logic EAL [7].

Changing Physical State. Assignment $x := \zeta$ performs instantaneous ontic change, e.g. pushing the autopilot button, *autopilot* := 1, or resetting a time counter with $t := 0$, as in Examples 1 and 2. No doxastic variables are allowed in ζ , since ontic truth is not directly a function of belief!

Differential equations $x' = \zeta \ \& \ \chi$ describe continuous motion over a nondeterministic duration, so long as the evolution domain constraint formula χ is true throughout. For example, $alt' = yv, t' = 1 \ \& \ t \leq 10$ describes linear change of altitude for up to 10 seconds according to vertical velocity yv . Nondeterministic ontic assignment $x := *$ is definable as $x' = 1; x' = -1$, which assigns any value in \mathbb{R} to x by increasing then decreasing it arbitrarily.

The test $?\phi$ transitions if and only if $d^4\mathcal{L}$ formula ϕ is true. It was used in Example 1 as an ontic trigger $?(alt > 1)$ determining whether an airplane could descend, and similarly as a belief trigger $?B_a(alt_a - \varepsilon > 1)$ in Example 2, where a pilot can only descend if they believe the airplane is safely above 1000 feet while taking worst-case noise into account.

Sequential composition $\alpha; \beta$ is self-explanatory. The choice $\alpha \cup \beta$ nondeterministically executes either α or β . It may be used to encode multiple possible outcomes or actions, e.g. $(?alt > 1; yv := -1) \cup yv := 1$ from Example 2.

Nondeterministic repetition α^* lets α be iterated arbitrarily many times. It was used in $(obs; ctrl; plant)^*$ to ensure the safety proof applies to a system that can run for a long time, not just to a one-time control decision.

Changing Belief State. Agent beliefs are updated by means of the *learning operator* $L_a(\gamma)$, where γ is a program describing belief change. Notably, to interleave ontic and belief change, the learning operator is a program itself rather than a modality as in [8,6]. Under $d^4\mathcal{L}$'s possible world semantics, each agent a considers multiple worlds possible. The intuitive behavior of $L_a(\gamma)$ is to execute program γ at each such world, and consider all outcomes of such executions as possible worlds.

The language of γ is a slightly modified subset of that of hybrid programs. Inside a learning operator, ontic assignment $x := \zeta$ becomes doxastic assignment $x_a := \theta$. Since doxastic change (unlike ontic change) may depend on previous beliefs, the assigned term θ allows doxastic variables. The language also includes test $?\phi$, choice $\gamma_1 \cup \gamma_2$ and sequential composition $\gamma_1; \gamma_2$.

This language of doxastic change captures the bulk of observation and reasoning phenomena found in belief-aware CPS, which tend to occur at distinct

and discrete intervals, e.g. looking at a sensor periodically. The literature [20,6] suggests that learned differential equations and repetition pose a very significant additional challenge, which is useful only in more specialized scenarios.

Learned programs may contain nondeterminism, as in $L_a(\gamma_1 \cup \gamma_2)$. Intuitively, this says that agent a is aware that either γ_1 or γ_2 happened, but cannot ascertain which: agent a must consider possible all outcomes of γ_1 and of γ_2 . Thus, in $\mathbf{d}^4\mathcal{L}$, learned nondeterminism is unobservable, and leads to the *indistinguishability of outcomes*, as in action models and epistemic actions [3,7]. This is contrast to program $L_a(\gamma_1) \cup L_a(\gamma_2)$, in which agent a either learns γ_1 , or learns γ_2 , but in both case knows precisely which one happened.

Learned test $L_a(? \xi)$ eliminates those possible worlds for which $? \xi$ does not succeed, i.e. in which ξ is false. In this way, $[L_a(? \phi)] \psi$ is analogous to public announcements and the tests of epistemic actions [7].

So far, the set of possible worlds may contract through learned tests and finitely expand with learned choice. The nondeterministic doxastic assignment $x_a := *$ further enables uncountable expansion of possibilities by assigning any value in \mathbb{R} to x_a . To let x_a take *any* value satisfying some property $\phi(x_a)$, the program $L_a(x_a := *; ? \phi(x_a))$ first “resets” the values x_a can take using nondeterministic assignment, and then contracts the set of possible worlds with $? \phi(x_a)$.

The grammar of programs divides programs into two categories. The first, denoted α , describes the language of ontic change, or the ontic fact $L_a(\gamma)$ that program γ was learned. The second, denoted γ , describes the language of doxastic change, and, as we have seen, is a subset of the first with minor modifications.

Definition 3. Let $x \in \mathbb{V}$, $a \in \mathbb{A}$, $x_a \in \mathbb{V}_a$ ϕ, ξ be formulas per Def. 2, θ, ζ be terms per Def. 1. Doxastic hybrid programs (DHP) α and learnable programs γ are defined thus:

$$\begin{array}{l} \alpha \quad ::= \quad x := \zeta \quad | \quad x' = \zeta \& \chi \quad | \quad ? \phi \quad | \quad \alpha; \alpha \quad | \quad \alpha \cup \alpha \quad | \quad \alpha^* \quad | \quad L_a(\gamma) \\ \gamma \quad ::= \quad x_a := \theta \quad | \quad x_a := * \quad | \quad ? \xi \quad | \quad \gamma; \gamma \quad | \quad \gamma \cup \gamma \end{array}$$

With a better understanding of $\mathbf{d}^4\mathcal{L}$ programs, we may now describe exactly how the belief of Example 3, $B_a((alt_a - alt)^2 \leq \varepsilon^2)$, is acquired.

Example 4 (Noisy sensors, cont'd). By observing a *trusted* altimeter, the pilot decides to forget previous beliefs about altitude and trust the current reading. Then, because the altimeter has a known error bound of ε , the pilot must now consider possible all altitude values at most ε away from the true value of alt .

$$L_a(alt_a := *; ?(alt_a - alt)^2 \leq \varepsilon^2)$$

4 Semantics of $\mathbf{d}^4\mathcal{L}$

The $\mathbf{d}^4\mathcal{L}$ semantics are designed to allow agents to hold potentially erroneous beliefs (proper belief, not knowledge) about a world which may undergo unobserved change. We are inspired by the modal Kripke semantics, but diverge from

it by completely decoupling the valuation describing ontic truth, denoted r in $d^4\mathcal{L}$, from agent beliefs, since unobservable actions must change ontic truth *only*.

Because beliefs are exclusively about the world and not about other beliefs, different agents' worlds need not interact with one another. Therefore, each agent a has their own set of worlds W_a , which they consider possible. Each agent a 's valuation $V_a(t)$ function holds the values of all doxastic variables at every world $t \in W_a$, e.g. agent a 's perception of altitude at $t \in W_a$ is $V_a(t)(alt_a)$.

In these sets of possible worlds, every world $t_1 \in W_a$ is indistinguishable from any other world $t_2 \in W_a$. Under the usual Kripke semantics, this means that the accessibility relation \sim_a , determining indistinguishability between worlds, is an equivalence relation, i.e. an S5 system. Equivalence relations traditionally encode knowledge, and belief is usually obtained by waiving the reflexivity requirement. In such belief systems, a *distinguished world* $s \in W_a$ determines ontic truth, and yet may not be accessible through \sim_a .

In $d^4\mathcal{L}$, we achieve belief by allowing discrepancies between the valuations of the possible worlds, including the distinguished one, and the separate ontic valuation r . Thus, a pilot could believe the airplane to be high with $V_a(t)(alt_a) > 1000$ for every $t \in W_a$, while it could be low in reality, with $r(alt) \leq 1000$.

This allows us to omit the accessibility relations entirely. It also simplifies learned program semantics since the learning operator can never inadvertently change ontic truth by altering the valuation of the distinguished world. We keep the distinguished world in Definition 4 as a means by which we may interpret every formula in every context, as we will see in Definitions 5 and 6.

This gives us the models of $d^4\mathcal{L}$, called physical-doxastic models, or PD-models for short. For simplicity, we consider only one agent a from now on, and we omit the subscript where it can be easily inferred, e.g. V instead of V_a .

Definition 4 (Physical/doxastic model). *A physical/doxastic model or PD-model $\omega = \langle r, W, V, s \rangle$ consists of 1) $r : \mathbb{V} \rightarrow \mathbb{R}$, the state of the physical world; 2) W a set of worlds called the possible worlds; 3) $V : W \rightarrow (\mathbb{V}_a \rightarrow \mathbb{R})$, a valuation function in which $V(t)(x_a)$ returns agent a 's perceived value of the doxastic variable x_a at world $t \in W$; and 4) $s \in W$, a distinguished world.*

PD-models are sufficient to give meaning to all terms, formulas and programs. We use ω, ν, μ to denote PD-models, and sub- and super-scripts are applied everywhere, e.g. $\omega' = \langle r', W', V', s' \rangle$. The shortcut $t \in \omega$ means $t \in W$; $\omega(t)(x_a)$ means $V(t)(x_a)$; and $\omega(x)$ means $r(x)$. The distinguished world of ω is $DW(\omega)$ and its distinguished valuation $DV(\omega) = \omega(DW(\omega)) = \omega(s) = V(s)$. The real world is $R(\omega) = r$. Finally, let $\langle r, W, V, s \rangle \oplus t = \langle r, W, V, t \rangle$ for any $t \in \omega$.

Interpretation of Terms, Formulas, and Programs. The interpretation of terms and formulas is standard, with logical variables X given meaning by a variable assignment $\eta : \Sigma \rightarrow \mathbb{R}$, state variables x by the physical state $R(\omega)$, and doxastic variables x_a by the distinguished valuation $DV(\omega)$. Terms and formulas such as alt_a and $alt_a > 1000$ may appear outside doxastic modalities during

calculus proofs. The distinguished valuation (for the distinguished world) ensures that they have a well-defined meaning and can thus be used as part of the proof.

Definition 5 (Term interpretation). *Let $\omega = \langle r, W, V, s \rangle$ be a PD-model, and $\eta : \Sigma \rightarrow \mathbb{R}$ be a logical variable assignment. Then, the interpretation of terms is defined inductively as follows: $val_\eta(\omega, x) = r(x)$ for state variable x ; $val_\eta(\omega, X) = \eta(X)$ for logical variable X ; $val_\eta(\omega, x_a) = DV(\omega)(x_a)$ for doxastic variable x_a ; $val_\eta(\omega, \theta_1 \otimes \theta_2) = val_\eta(\omega, \theta_1) \otimes val_\eta(\omega, \theta_2)$ for $\otimes \in \{+, -, \times, \div\}$.*

Formula interpretation is derived directly from $d\mathcal{L}$, first-order logic for real arithmetic, and simplified Kripke semantics for beliefs. Definitions 6 and 7 are mutually recursive due to the box modality formula $[\alpha]\phi$ and test program $?\phi$.

Definition 6 (Interpretation of formulas). *Let $\omega = \langle r, W, V, s \rangle$ be a PD-model, η be a variable assignment, and $\langle r, W, V, s \rangle \oplus t = \langle r, W, V, t \rangle$. Then, the valuation of a formula ϕ as 1 (true) or 0 (false) is defined inductively as follows.*

$$\begin{array}{ll}
 val_\eta(\omega, \theta_1 \leq \theta_2) = 1 & \text{iff } val_\eta(\omega, \theta_1) \leq val_\eta(\omega, \theta_2) \\
 val_\eta(\omega, \phi_1 \vee \phi_2) = 1 & \text{iff } val_\eta(\omega, \phi_1) = 1 \text{ or } val_\eta(\omega, \phi_2) = 1 \\
 val_\eta(\omega, \neg\phi) = 1 & \text{iff } val_\eta(\omega, \phi) = 0 \\
 val_\eta(\omega, \forall X \phi) = 1 & \text{iff for all } v \in \mathbb{R}, val_{\eta[X \mapsto v]}(\omega, \phi) = 1 \\
 val_\eta(\omega, B_a(\xi)) = 1 & \text{iff for all } t \in \omega, val_\eta(\omega \oplus t, \xi) = 1 \\
 val_\eta(\omega, [\alpha]\phi) = 1 & \text{iff for all } (\omega, \omega') \in \rho_\eta(\alpha), val_\eta(\omega', \phi) = 1
 \end{array}$$

Under these semantics, $B_a(x = 0)$ is equivalent to $x = 0$ since state variable x is independent of the choice of distinguished world, unlike x_a . CPS designers have no reason to write such formulas, but when they do appear in calculus proofs, the doxastic modality is eliminated using the equivalence $B_a(x = 0) \leftrightarrow x = 0$.

Program Semantics The program semantics is given as a reachability relation over PD-models, with $(\omega, \omega') \in \rho_\eta(\alpha)$ meaning that PD-model ω' is reachable from ω using program α . The semantics of DHPs starts with that of $d\mathcal{L}$'s hybrid programs. Most cases are intuitive. Differential equations use their solution y to evolve $R(\omega)$ for a nondeterministic duration, and ensure the evolution domain constraint χ is satisfied throughout. For a more in-depth treatment, see [17].

To this we add doxastic assignment, which affects the distinguished valuation $DV(\omega)$, and the learning operator, which represents the “execute γ at each possible world” semantics from DELs, as illustrated in Figure 1.

In Figure 1, let $(\omega, \omega') \in \rho_\eta(L_a(\gamma))$. Then, each world $\nu \in \omega'$ after learning has an “origin” world $t \in \omega$ from before learning, e.g. t_1 is the origin world for ν_1 and ν_2 . Every PD-model ν that γ can reach from each origin world $t \in \omega$ (i.e. $(\omega \oplus t, \nu) \in \rho_\eta(\gamma)$) becomes a possible world $\nu \in \omega'$ after $L_a(\gamma)$. The valuation $\omega'(\nu)$ reflects the effects of γ , which can be found in the distinguished valuation of ν , and thus, we let $\omega'(\nu) = DV(\nu)$.

Finally, the distinguished world of ω' is chosen as any $t' \in \omega'$ whose origin world is $\text{DW}(\omega)$. This applies the principle of *learned nondeterminism as indistinguishability of outcomes to the distinguished world*.




Fig. 1. The double-circled $t_1 = \text{DW}(\omega)$ creates, through γ 's nondeterminism, two post-learning worlds $\nu_1, \nu_2 \in \omega'$ worlds, either of which can be nondeterministically chosen as $\text{DW}(\omega')$. The world $t_2 \in \omega$ leads to $\nu_3 \in \omega'$, which cannot be chosen as $\text{DW}(\omega')$.

Definition 7 (Transition semantics). Let $\omega = \langle r, W, V, s \rangle$ be a PD-model, and η be a variable assignment. The transition relation for doxastic dynamic programs is inductively defined by:

- $(\omega, \omega') \in \rho_\eta(x := \zeta)$ iff $\omega' = \omega$ except $R(\omega')(x) = \text{val}_\eta(\omega, \zeta)$
- $(\omega, \omega') \in \rho_\eta(x_a := \theta)$ iff $\omega' = \omega$ except $DV(\omega')(x_a) = \text{val}_\eta(\omega, \theta)$
- $(\omega, \omega') \in \rho_\eta(x_a := *)$ iff $\omega' = \omega$ except $DV(\omega')(x_a) = v$ for some $v \in \mathbb{R}$
- $(\omega, \omega') \in \rho_\eta(x' = \zeta \ \& \ \chi)$ iff $\omega' = \langle r[x \mapsto y(\tau)], W, V, s \rangle$ for the solution $y : [0, T] \rightarrow \mathbb{R}$ of the diff. eq., with $\tau \in [0, T]$ for some $T \geq 0$. Furthermore, for all $t_i \in [0, \tau]$, and $\text{val}_\eta(\langle r[x \mapsto y(t_i)], W, V, s \rangle, \chi) = 1$.
- $(\omega, \omega) \in \rho_\eta(? \phi)$ iff $\text{val}_\eta(\omega, \phi) = 1$
- $\rho_\eta(\alpha; \beta) = \rho_\eta(\alpha) \circ \rho_\eta(\beta)$
 $= \{\omega_3 : \text{there is } \omega_2 \text{ s.t. } (\omega_1, \omega_2) \in \rho_\eta(\alpha) \text{ and } (\omega_2, \omega_3) \in \rho_\eta(\beta)\}$
- $\rho_\eta(\alpha \cup \beta) = \rho_\eta(\alpha) \cup \rho_\eta(\beta)$
- $(\omega, \omega') \in \rho_\eta(\alpha^*)$ iff there is $n \in \mathbb{N}$ such that $(\omega, \omega') \in \rho_\eta(\alpha^n)$, where α^n is α sequentially composed n times.
- $(\omega, \omega') \in \rho_\eta(L(\gamma))$ if: $r' = r$, $W' = \{\nu : \text{there is } t \in \omega \text{ s.t. } (\omega \oplus t, \nu) \in \rho_\eta(\gamma)\}$, $\omega'(\nu) = DV(\nu)$ for all $\nu \in \omega'$, and $\text{DW}(\text{DW}(\omega')) = \text{DW}(\omega)$.

Figure 1 and Definition 7 show that $d^4\mathcal{L}$'s learning operator applies the DEL semantics to *any* language of change, so long as it has a transition semantics, as in $(\omega \oplus t, \nu) \in \rho_\eta(\gamma)$. It is possible to extend this operator to traditional multi-agent Kripke structures by letting two after-learning worlds be indistinguishable in ω' iff their origin worlds were indistinguishable in ω , as is standard in DELs.

5 Sound Sequent Calculus

Our main contribution towards the verification of belief-aware CPS is a sound proof calculus for $d^4\mathcal{L}$. The meaning of a sequent $\Gamma \vdash \phi$ with a $d^4\mathcal{L}$ formula ϕ and a set of $d^4\mathcal{L}$ formulas Γ is captured with the following definition of validity.

Definition 8 (Validity). *A sequent $\Gamma \vdash \phi$ is valid iff for all ω and η ,*

$$val_\eta(\omega, \bigwedge_{\psi \in \Gamma} \psi \rightarrow \phi) = 1$$

For simplicity's sake, we use a single definition of soundness for proof rules.

Definition 9 (Global Soundness). *A proof rule PR , as in $PR \frac{\Gamma_1 \vdash \phi_1}{\Gamma_2 \vdash \phi_2}$, is globally sound when, if $\Gamma_1 \vdash \phi_1$ is valid then $\Gamma_2 \vdash \phi_2$ is valid.*

Overview of the Calculus. Figure 2 contains the fragment of the calculus that pertains to the learning operator. The $d\mathcal{L}$ calculus [16] is omitted as it is easily adaptable to $d^4\mathcal{L}$. Single-modality agent rationality axioms can be adopted for belief, i.e. $B_a(\phi_1 \rightarrow \phi_2) \rightarrow (B_a(\phi_1) \rightarrow B_a(\phi_2))$ and, if ϕ is valid, then $B_a(\phi)$ is too. The proof for the following theorem can be found in [13].

Theorem 1. *The proof rules in Figure 2 are globally sound.*

Sequent contexts Γ are partitioned into $\Gamma_R; \Gamma_B; \Gamma_P; \Gamma_O$. The set Γ_R is the set of formulas with only state and logical variables and no doxastic modalities, e.g. $alt > 0$. Γ_B and Γ_P are the sets of belief and possibility formulas respectively, e.g. $B_a((alt_a - alt)^2 \leq \varepsilon^2)$ and $P_a((alt_a - alt)^2 \leq \varepsilon^2)$. Γ_O is the set of formulas with doxastic variables but no modalities, e.g. $alt_a > 0$. The rules in Figure 2 are only applicable once this partitioning has been achieved. Finally, let $\Gamma \setminus_{x_a} = \{\phi \in \Gamma : x_a \text{ does not occur in } \phi\}$.

Proof rules for learned programs that change doxastic state, like assignment or test, work by altering the contexts in suitable ways. Each learned program has two rules, for the \square and \diamond dynamic modalities, which deal with the non-determinism in the choice of the distinguished world. The exception is $L_a(\alpha \cup \beta)$, where doxastic and dynamic modalities interact much more subtly.

The proof rules for assignment $L_a(x_a := \theta)$ capture the intuition that, since x_a now has the value of θ at each possible world, syntactically substituting all occurrences of x_a with θ ought to mean the same thing.

Since nondeterministic assignment $L_a(x_a := *)$ gives x_a any possible value, then anything previously possible about x_a remains possible. However, beliefs about x_a , which must hold for *all* worlds, do not survive the assignment (unless they are tautologies). The proof rules $[L := *]$ and $\langle L := * \rangle$ eliminate the formulas which may no longer hold after assignment from the context.

Formulas describing the distinguished world, i.e. in Γ_O , are retained or removed, respectively, depending on whether the dynamic modality allows us pick our distinguished world to suit our goals, as with \diamond , or not, as with \square .

$$\begin{array}{l}
[L:=] \frac{\Gamma \vdash \phi(\theta)}{\Gamma \vdash [L_a(x_a := \theta)] \phi(x_a)}^1 \\
[L:=*] \frac{\Gamma_R; \Gamma_B \setminus_{x_a}; \Gamma_P; \Gamma_O \setminus_{x_a} \vdash \phi}{\Gamma \vdash [L_a(x_a := *)] \phi} \\
[L?] \frac{\Gamma_R; \Gamma_B; \emptyset; \Gamma_O \vdash B_a(\xi) \rightarrow \psi}{\Gamma_R; \Gamma_B; \Gamma_P; \Gamma_O \vdash [L_a(?\xi)] \psi} \\
[L:] \frac{\Gamma \vdash [L_a(\gamma_1); L_a(\gamma_2)] \phi}{\Gamma \vdash [L_a(\gamma_1; \gamma_2)] \phi} \\
[LB\cup] \frac{\Gamma \vdash [L_a(\gamma_1)] B_a(\xi) \wedge [L_a(\gamma_2)] B_a(\xi)}{\Gamma \vdash [L_a(\gamma_1 \cup \gamma_2)] B_a(\xi)} \\
[LPU] \frac{\Gamma \vdash [L_a(\gamma_1)] P_a(\xi) \wedge [L_a(\gamma_2)] P_a(\xi)}{\Gamma \vdash [L_a(\gamma_1 \cup \gamma_2)] P_a(\xi)} \\
[LU] \frac{\Gamma \vdash [L_a(\gamma_1)] \phi \wedge [L_a(\gamma_2)] \phi}{\Gamma \vdash [L_a(\gamma_1 \cup \gamma_2)] \phi}^2 \\
\langle L:= \rangle \frac{\Gamma \vdash \phi(\theta)}{\Gamma \vdash \langle L_a(x_a := \theta) \rangle \phi(x_a)}^1 \\
\langle L:=* \rangle \frac{\Gamma_R; \Gamma_B \setminus_{x_a}; \Gamma_P; \Gamma_O \vdash \phi}{\Gamma \vdash \langle L_a(x_a := *) \rangle \phi} \\
\langle L? \rangle \frac{\Gamma_R; \Gamma_B; \emptyset; \Gamma_O \vdash B_a(\xi) \wedge \psi}{\Gamma_R; \Gamma_B; \Gamma_P; \Gamma_O \vdash \langle L_a(?\xi) \rangle \psi} \\
\langle L; \rangle \frac{\Gamma \vdash \langle L_a(\gamma_1); L_a(\gamma_2) \rangle \phi}{\Gamma \vdash \langle L_a(\gamma_1; \gamma_2) \rangle \phi} \\
\langle LB\cup \rangle \frac{\Gamma \vdash \langle L_a(\gamma_1) \rangle B_a(\xi) \wedge \langle L_a(\gamma_2) \rangle B_a(\xi)}{\Gamma \vdash \langle L_a(\gamma_1 \cup \gamma_2) \rangle B_a(\xi)} \\
\langle LPU \rangle \frac{\Gamma \vdash \langle L_a(\gamma_1) \rangle P_a(\xi) \vee \langle L_a(\gamma_2) \rangle P_a(\xi)}{\Gamma \vdash \langle L_a(\gamma_1 \cup \gamma_2) \rangle P_a(\xi)} \\
\langle LU \rangle \frac{\Gamma \vdash \langle L_a(\gamma_1) \rangle \phi \vee \langle L_a(\gamma_2) \rangle \phi}{\Gamma \vdash \langle L_a(\gamma_1 \cup \gamma_2) \rangle \phi}^2
\end{array}$$

¹ The substitution of x_a by θ must be admissible in ϕ , see **Doxastic Assignment**

² Formula ϕ does not contain doxastic modalities or variables, or learning operators

Fig. 2. Dynamic doxastic fragment of the $d^4\mathcal{L}$ calculus, with Γ being $\Gamma_R; \Gamma_B; \Gamma_P; \Gamma_O$

Learned test results in the belief about the test result, as in public announcements. The test contracts the set of possible worlds, so we must remove the set of possibility formulas from the context, as they may no longer hold. The underlying dynamic modality determines whether this belief is a precondition for ψ or a necessity (\diamond implies at least one transition, \square does not).

Learned sequential composition is merely reduced to regular sequential composition. Doxastic assignment and choice deserve further attention below.

Doxastic Assignment. The rule for doxastic assignments relies on its syntactic substitution being equivalent to the semantic substitution effected by learned assignment. This nontrivial result can be captured succinctly by Lemma 1, whose full proof is found in [13]. This result only holds when the substitution is *admissible* with respect to a given formula ϕ , i.e. that syntactic conditions are in place ensuring the substitution will not change the meaning of the substituted variables, and therefore, of the formula [13].

Lemma 1 (Doxastic Substitution Lemma). *Let ϕ be a formula. Let σ be an admissible substitution for ϕ which replaces only doxastic variable x_a . Then, for every η and $\omega = \langle r, W, V, s \rangle$, we have $val_\eta(\omega, \sigma(\phi)) = val_\eta(\sigma(\omega), \phi)$, where $\sigma(\phi)$ is syntactic substitution, and $\sigma(\omega)$ is semantic substitution, defined as $\sigma(\omega) = \langle r, W, \sigma(V), s \rangle$, with $\sigma(V)(t)(x_a) = val_\eta(\omega \oplus t, \sigma(x_a))$ and $\sigma(V)(t)(y_a) = V(t)(y_a) = \omega(t)(y_a)$ for $y_a \neq x_a$, for all $t \in \omega$.*

Nondeterministic Choice. Learned choice influences *doxastic* modalities, and the choice of distinguished world is influenced by *dynamic* modalities. This makes

for some subtlety in the rules for learned choice. Consider the potential rule below, which assumes $L_a(\gamma_1 \cup \gamma_2)$ is equivalent to $L_a(\gamma_1) \cup L_a(\gamma_2)$.

$$\frac{P_a(\neg\xi); \xi \vdash \langle L_a(?\xi) \rangle B_a(\xi) \vee \langle L_a(?True) \rangle B_a(\xi)}{P_a(\neg\xi); \xi \vdash \langle L_a(? \xi \cup ?True) \rangle B_a(\xi)}$$

The sequent contexts tell us that ξ holds in the distinguished world $\text{DW}(\omega)$, but not in some other $t \in \omega$. The disjunction holds, since $\langle L_a(? \xi) \rangle B_a(\xi)$ is trivially true. The program $L_a(? \xi \cup ?True)$ preserves all worlds, including t , because of $?True$. Since ξ is not true in t , agent a cannot therefore believe ξ . But if the top is valid and the bottom is not, this rule would be unsound.

This phenomenon occurs because the conclusion of the rule requires us to prove $B_a(\xi)$ for worlds originated through both $? \xi$ and $?True$. However, the premise of the rule implies we need only check those from either $? \xi$ or $?True$, as if the \diamond dynamic modality had control over learned nondeterminism. It does not: outcomes of learned nondeterminism are *always* considered indistinguishable.

Proof rules for $L_a(\gamma_1 \cup \gamma_2)$ must therefore be as conservative as the most conservative of their dynamic and doxastic modalities: the only proof rule that allows disjunction in the premise is $\langle LP \cup \rangle$ since both modalities \diamond and $P_a(\cdot)$ are existential. This realization informs the soundness proofs for learned choice.

Proof (Soundness sketch for $\langle LB \cup \rangle$). Let ω be an arbitrary PD-model. We must show that $\text{val}_\eta(\omega, \langle L_a(\gamma_1 \cup \gamma_2) \rangle B_a(\xi)) = 1$, i.e. that ξ is true at every world ν reachable by either $(t, \nu) \in \rho_\eta(\gamma_1)$ or $(t, \nu) \in \rho_\eta(\gamma_2)$ for $t \in \omega$.

Let $(t, \nu) \in \rho_\eta(\gamma_1)$. By hypothesis, $\text{val}_\eta(\omega, \langle L_a(\gamma_1) \rangle B_a(\xi)) = 1$, i.e. ξ is true at every world reachable by γ_1 , and ν in particular. The argument is symmetrical for $(t, \nu) \in \rho_\eta(\gamma_2)$, but only because the premise is a conjunction. Thus, for any world ν created by $L_a(\gamma_1 \cup \gamma_2)$, ξ is true at that world. Therefore, $B_a(\xi)$. \square

6 Validation and Application

We will now use $\text{d}^4\mathcal{L}$ to illustrate how to prove the safety of a small belief-aware CPS. The scenario is similar to that of Example 2, and it is useful to have a reference for some of the most used $\text{d}\mathcal{L}$ proof rules that $\text{d}^4\mathcal{L}$ inherits [16].

$$[;] \frac{\Gamma \vdash [\alpha] [\beta] \phi}{\Gamma \vdash [\alpha; \beta] \phi} \quad [?] \frac{\Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash [?\phi] \psi} \quad \rightarrow R \frac{\Gamma, \phi \vdash \psi}{\Gamma \vdash \phi \rightarrow \psi}$$

We let the pilot observe the altimeter with $\mathbf{0} \equiv L_a(\text{alt}_a := *; ?Noise)$, with $Noise \equiv (\text{alt}_a - \text{alt} < \varepsilon)$. The control program \mathbf{C} climbs or descends by setting vertical velocity depending on whether descent is believed to be safe, $\mathbf{CB} \cup \mathbf{CP} \equiv (?B_a(\text{alt}_a - T - \varepsilon > 0); yv := -1) \cup (?P_a(\text{alt}_a - T - \varepsilon \leq 0); yv := 1)$. The two tests are mutually exclusive, leading to dual belief operators: descending requires the strong condition of belief, whereas the mere possibility of being too low triggers a climb. We use $\mathbf{F} \equiv t := 0; t' = 1, \text{alt}' = yv \ \& \ t < T$ as very simplified flight dynamics, and an invariant $\text{inv} \equiv (\text{alt} > 0 \wedge T > 0)$ to handle repetition.

We will prove the validity of the formula $alt > 0, T > 0 \vdash [(0; C; F)^*] alt > 0$ by successively applying sound proof rules from $d\mathcal{L}$ and Figure 2 to it. The leaves of the proof tree will be formulas that can be easily discharged using only $d\mathcal{L}$ rules or real arithmetic. Once the proof tree is complete, we will know this safety formula is valid, and thus that the modeled system is safe.

$$\text{loop} \frac{\frac{*}{alt > 0, T > 0 \vdash inv} \quad [;] [;] \frac{inv \vdash [0] [C] [F] inv}{inv \vdash [0; C; F] inv} \quad \frac{*}{inv \vdash alt > 0}}{alt > 0, T > 0 \vdash [(0; C; F)^*] alt > 0}$$

The middle branch continues in:

$$[\cup] \frac{inv; B_a(Noise) \vdash [CB] [F] inv \quad inv; B_a(Noise) \vdash [CP] [F] inv}{\frac{[L?] \rightarrow R \frac{inv; B_a(Noise) \vdash [C] [F] inv}{inv \vdash [L_a(?Noise)] [C] [F] inv}}{[L:=*] \frac{inv \vdash [L_a(alt_a := *)] [L_a(?Noise)] [C] [F] inv}{[L;] [;] \frac{inv \vdash [L_a(alt_a := *)] [L_a(?Noise)] [C] [F] inv}}{inv \vdash [L_a(alt_a := *)] [C] [F] inv}}}$$

The branch on the right closes using $d\mathcal{L}$ proof rules and standard $d\mathcal{L}$ reasoning independent of beliefs: if the airplane is above ground and climbs, it remains above ground. The left branch requires some doxastic reasoning.

$$\text{cut} \frac{inv; B_a(Noise), B_a(alt_a - T - \varepsilon > 0) \vdash alt > T \quad inv; alt > T \vdash [F(-1)] inv}{\frac{[:=] \frac{inv; B_a(Noise), B_a(alt_a - T - \varepsilon > 0) \vdash [F(-1)] inv}{[;] [?] \rightarrow R \frac{inv; B_a(Noise), B_a(alt_a - T - \varepsilon > 0) \vdash [yv := -1] [F(yv)] inv}{inv; B_a(Noise) \vdash [?B_a(alt_a - T - \varepsilon > 0); yv := -1] [F(yv)] inv}}}$$

The left side of the cut rule must show that $alt > T$, and for that we will use the S5 rationality axioms that allow for reasoning about arithmetic. Thus, the agent may conclude (1) $B_a(alt > alt_a - \varepsilon)$ from $B_a(Noise)$, and (2) $B_a(alt_a > T + \varepsilon)$ from $B_a(alt_a - T - \varepsilon > 0)$. But (1) and (2) together lead to $B_a(alt > T)$, which no longer contains any doxastic variables. It is therefore equivalent to $alt > T$. We have thus used the belief meta-property (1), relating ontic and doxastic truth, to obtain an important fact about the world which we may now use in the right side of the proof.

This right side is a standard $d\mathcal{L}$ proof without doxastics: the rules for differential equations show that, after evolving for at most T time at a speed of -1 , the airplane cannot end up below ground, since it started above T altitude.

This completes the sequent proof. It leveraged a mix of ontic, doxastic and meta-doxastic statements in order to make the argument for the safety of this controller. When working with trusted sensors, we also see an intuitive partitioning of the proof: first, doxastic formulas ($B_a(alt_a - T - \varepsilon > 0)$) and meta-doxastic formulas ($B_a(Noise)$) are used to derive ontic formulas ($alt > T$). Second, such ontic statements form the basis for arguments made in $d\mathcal{L}$ -exclusive proof branches that ensure post-control actuation results in safe behavior. This clear separation of concerns allows CPS engineers to work more intuitively and compositionally during the design and verification stages of belief-aware CPS.

The ways in which agents learn and reason influence the ontic facts that can be deduced, but those facts must in turn be informed by safety requirements of the CPS's physical evolution. Doxastics and ontics clearly play off each and have, in the past, contributed to safety incidents. By making this explicit in the model, $d^4\mathcal{L}$ ensures adequate attention is given to such dynamics so that hopefully, ontic/doxastic concerns can be identified before they lead to tragedy.

7 Related Work

The logic $d^4\mathcal{L}$ takes heavy inspiration from two bodies of work: one for reasoning about a changing world, and one for reasoning about changing beliefs.

Changing world The logic $d\mathcal{L}$ for reasoning about the ontic dynamics of CPS [16,17,19] has shown itself to be capable of verifying interesting and relevant real world systems [18,17]. However, it requires manual modeling discipline to express noise [14], rather than having noise or beliefs thereof as built-in primitives.

The example used in this paper is so simple that it can still be converted to $d\mathcal{L}$ using modeling tricks [14]. The trick is to transform alt_a into a state variable and remove the learning operator from the observation program, i.e. $alt_a := *; ?Noise$ rather than $L_a(alt_a := *; ?Noise)$. The agent's control would then be $(?alt_a - T - \varepsilon > 0; yv := -1) \cup (?alt_a - T - \varepsilon \leq 0; yv := 1)$.

However, this conversion relies fundamentally on the box dynamic modality $[\alpha]\phi$, which checks safety for *all* executions of $alt_a := *; ?Noise$. With liveness formulas using the diamond dynamic modality $\langle \alpha \rangle \phi$, safety need only be checked for *one* execution. Thus, in liveness formulas, this method would fail to capture the intended behavior of both the learning operator and the belief modality, which should still apply to *all* possible worlds, or, in $d\mathcal{L}$ terms, all executions.

This conversion can also quickly become complex. A more detailed controller for a pilot trying to remain around or above cruising altitude A could be $(?B_a(alt_a - T - \varepsilon > A); yv := -1) \cup (?P_a(alt_a - T - \varepsilon > A); yv := -0.5) \cup (?B_a(alt_a - T - \varepsilon \leq 0); yv := 1)$. This is similar to previous controllers, but allows for a more gentle descent when the pilot considers the possibility of being close to A . The equivalent $d\mathcal{L}$ controller is $(?alt_a - T - \varepsilon > A; yv := -1) \cup (?alt_a - T + \varepsilon > A; yv := -0.5) \cup (?alt_a - T - \varepsilon \leq 0; yv := 1)$. However, this elimination of doxastic modalities requires a change in the arithmetic itself, e.g. $(?P_a(alt_a - T - \varepsilon > A))$ turns into $(?alt_a - T + \varepsilon > A)$. Belief must consider worst case noise, whereas possibility can consider the best case. This can quickly become complex when going beyond simpler interval-based noise scenarios.

Both $d\mathcal{L}$ and $d^4\mathcal{L}$ controllers allow tests for deciding which action to take, but represent action triggers in first-order logic or doxastic logic, respectively, e.g. $alt_a - T + \varepsilon > A$ and $P_a(alt_a - T - \varepsilon > A)$. Decisions in real CPS are based on belief, and as the conversion from doxastic to non-doxastic action triggers quickly becomes non-trivial, it is best to avoid subtle modeling mistakes by working with belief during design and verification. With $d^4\mathcal{L}$, safety engineers can rely on doxastic intuitions during verification, rather than having to infer

them from formulas such as $alt_a - T + \varepsilon > A$, which does not clearly convey the concept of possibility that is so clear in $P_a(alt_a - T - \varepsilon > A)$.

The notion of robustness in hybrid systems control can capture complex notions of sensor and actuator noise [11], but is ultimately restrictive for the purpose of belief-aware CPS, as discussed at the beginning of Section 2. Adaptive control, where no *a priori* constraints are known, often depends on neural networks [15], and safety guarantees for systems relying on learning are known to add significant complexity to such efforts [9].

Changing belief On the other side, we have dynamic epistemic logics (DELS) [6,3,4,10,7], of which a good overview can be found in the literature [8]. They provide several notions of learning for different languages, some similar to our programs [6]. *Public* propositional world-change [6] would make ontic change implicitly observable, which is in direct conflict with the unobservability requirements of belief-aware CPS. Furthermore, relevant DEL axiomatizations rely on creating a conjunction out of properties of each accessible possible world [4,8], which is incompatible with the uncountably many worlds that CPS demand.

Belief revision through the AGM postulates [2] is an axiomatic, declarative approach to belief change. Because it is such a different approach, it presents many challenges in its integration with model-theoretic work such as $d\mathcal{L}$.

In order to begin addressing safety concerns around ontic/doxastic interactions at design time, CPS engineers and agents must make complex logical arguments from both ontic facts and beliefs, as in Section 6. Despite their many successes, the work describe in this section do not address this particular challenge directly in a principled way.

8 Conclusions

This paper considers interactions between belief and fact, which have significant safety implications. We proposed belief-aware CPSs as a first-principles paradigm under which safety concerns with such ontic/doxastic dynamics are expressly dealt with at design time, before safety violations occur. Our contribution is the logic $d^4\mathcal{L}$ for modeling and verifying belief-aware CPSs, requiring simultaneous, complex belief- and world-change. Its formulas can describe ontic, doxastic and meta-doxastic statements, and its programs can model belief-aware CPS with belief-triggered controllers that make decisions based only on what they can observe and reason. We proposed a learning operator for belief-change, which is capable of transforming any transition-based semantics of change into a semantics of *belief-change*. We presented a sequent calculus for $d^4\mathcal{L}$, which is proven to be sound, and used it to show the safety of a simple belief-aware CPS. This is, to the best of our knowledge, the first calculus for a dynamic logic of belief/knowledge change that can handle an uncountable domain, as in CPS.

Acknowledgment. We thank the anonymous reviewers for their helpful feedback.

References

1. Aircraft Accident Investigation Bureau of Ethiopia: Report No. AI-01/19, Aircraft Accident Investigation Preliminary Report, Ethiopian Airlines Group, B737-8 (MAX) Registered ET-AVJ (2019)
2. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.* **50**(2), 510–530 (1985). <https://doi.org/10.2307/2274239>
3. Baltag, A., Moss, L.S.: Logics for epistemic programs. *Synthese* **139**(2), 165–224 (2004). <https://doi.org/10.1023/B:SYNT.0000024912.56773.5e>
4. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: TARK. pp. 43–56. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1998), <http://dl.acm.org/citation.cfm?id=645876.671885>
5. Bureau d’Enquêtes et d’Analyses (BEA): Final report on the accident on 1st june 2009 to the airbus A330-203 registered F-GZCP operated by Air France flight AF 447 from Rio de Janeiro to Paris (2012)
6. van Ditmarsch, H.P., van der Hoek, W., Kooi, B.P.: Dynamic epistemic logic with assignment. In: AAMAS. pp. 141–148. ACM, New York, NY, USA (2005). <https://doi.org/10.1145/1082473.1082495>, <http://doi.acm.org/10.1145/1082473.1082495>
7. van Ditmarsch, H.P.: Descriptions of game actions. *Journal of Logic, Language and Information* **11**(3), 349–365 (2002). <https://doi.org/10.1023/A:1015590229647>
8. Ditmarsch, H.v., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*. Springer (2007)
9. Fulton, N., Platzer, A.: Verifiably safe off-model reinforcement learning. In: Vojnar, T., Zhang, L. (eds.) TACAS. LNCS, vol. 11427, pp. 413–430. Springer (2019). https://doi.org/10.1007/978-3-030-17462-0_28
10. Gerbrandy, J., Groeneveld, W.: Reasoning about information change. *Journal of Logic, Language and Information* **6**(2), 147–169 (1997). <https://doi.org/10.1023/A:1008222603071>
11. Goebel, R., Hespanha, J.P., Teel, A.R., Cai, C., Sanfelice, R.: Hybrid systems: Generalized solutions and robust stability. In: Proc. of the 6th IFAC Symp. on Nonlinear Contr. Systems (Sep 2004)
12. Komite Nasional Keselamatan Transportasi: Preliminary Aircraft Accident Investigation Report, PT. Lion Mentari Airlines, Boeing 737-8 (MAX); PK-LQP (2018)
13. Martins, J.G., Platzer, A., Leite, J.: A sound calculus for a logic of belief-aware cyber-physical systems. Tech. Rep. CMU-CS-19-116, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA (July 2019)
14. Mitsch, S., Ghorbal, K., Vogelbacher, D., Platzer, A.: Formal verification of obstacle avoidance and navigation of ground robots. I. *J. Robotics Res.* **36**(12), 1312–1340 (2017). <https://doi.org/10.1177/0278364917733549>
15. Nguyen, N.T., Krishnakumar, K.S., Kaneshige, J.T., Nespeca, P.P.: Flight dynamics and hybrid adaptive control of damaged aircraft. *Journal of Guidance, Control, and Dynamics* **31**(3), 751–764 (2008). <https://doi.org/10.2514/1.28142>
16. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reas.* **41**(2), 143–189 (2008). <https://doi.org/10.1007/s10817-008-9103-8>
17. Platzer, A.: Logics of dynamical systems. In: LICS. pp. 13–24. IEEE (2012). <https://doi.org/10.1109/LICS.2012.13>

18. Platzer, A.: Logic & proofs for cyber-physical systems. In: Olivetti, N., Tiwari, A. (eds.) IJCAR. LNCS, vol. 9706, pp. 15–21. Springer (2016). https://doi.org/10.1007/978-3-319-40229-1_3
19. Platzer, A.: Logical Foundations of Cyber-Physical Systems. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-63588-0>
20. Platzer, A., Tan, Y.K.: Differential equation axiomatization: The impressive power of differential ghosts. In: Dawar, A., Grädel, E. (eds.) LICS. pp. 819–828. ACM, New York (2018). <https://doi.org/10.1145/3209108.3209147>