

Extending Dynamic Logics with First-Class Relational Reasoning

Jian Xiang¹ and Stephen Chong²

¹ UNC Charlotte, 8723 Cameron Blvd, Charlotte, NC 28262, USA

² Harvard University, 150 Western Ave, Boston, MA 02134, USA

Abstract. Many important properties of computer systems are relational properties, which are often difficult to express and verify. Dynamic logics are well-known formalisms for program verification. We present a general extension, called the REL extension, of dynamic logics to support first-class relational reasoning. The extension provides intuitive syntax to express relational properties, which may be difficult or impossible to express in the host dynamic logic. The REL extension can be instantiated for different host logics. Verifying relational properties expressed by the REL extension can benefit from techniques developed for general relational reasoning and domain-specific relational reasoning, and existing tools developed for the host logic.

We validate the applicability of the REL extension by instantiating it for a well-known dynamic logic: differential dynamic logic ($\text{d}\mathcal{L}$). As a result, the instantiation can express key relational properties that cannot be easily expressed with $\text{d}\mathcal{L}$. We develop an encoding for the instantiation to leverage existing verification tools for $\text{d}\mathcal{L}$. We conduct an experiment on a set of benchmarks, and successfully verify a set of non-trivial relational properties either fully or semi automatically.

Keywords: relational reasoning · dynamic logic · formal verification.

1 Introduction

Many important properties, such as noninterference [39], are relational properties [18], i.e., predicates over sets of pairs of executions. They are fundamental properties but often harder to express and reason about than single-trace properties, as they require reasoning simultaneously about multiple executions. Moreover, there is less tool support for verification of relational properties, compared to single-trace properties.

Dynamic logics are multi-modal logics widely used for verifying single-trace properties of imperative programs [40, 1, 34, 28]. They are specified over a set of programs α and a set of formulas ϕ . Program specifications are stated with modality of necessity $[\alpha]\phi$ that reads “after any execution of α , ϕ is true” and modality of existence $\langle\alpha\rangle\phi$ that reads “after some execution of α , ϕ is true”. Different variants of dynamic logics have been introduced and used in different application domains. For example, differential dynamic logic is developed for the

safety of cyber-physical systems [37, 35], and linear dynamic logic is often used for AI-based planning [20]. Various tools and techniques have been developed for the verification of formulas expressed with these dynamic logics. However, relational verification is often not well-explored across many of these application domains. It would be beneficial if we can incorporate relational reasoning into these dynamic logics and leverage existing tools for relational verification.

A lightweight yet efficient approach for expressing relational properties is extending the host logic with a *biprogram* construct [38, 6], which explicitly specifies a pair of programs and thus allows reasoning about executions of the pair. We argue that extending dynamic logics with a biprogram provides intuitive and effective support for expressing relational properties. Elements from both programs can be directly used in expressing relational properties. Modalities can be designed to quantify over executions of the program pair. As such, a dynamic logic can be cleanly extended with first-class relational reasoning.

We contribute a general extension, called the REL extension, that builds upon the biprogram construct to support first-class relational reasoning for different dynamic logics. A REL extension has three major components³:

- REL programs that specify pairs of programs. A *biprogram* (α, β) can specify a *left* program (i.e., α) and a *right* program (i.e., β), of the host logic.
- REL modalities that express executions of program pairs. They play the same role as modalities in dynamic logics: express different quantifications over executions of program pairs. For example, REL modality of necessity $\llbracket (\alpha, \beta) \rrbracket \phi_{\tau}$ expresses that *for all* executions of α and β , ϕ_{τ} holds at their last states.
- REL formulas that directly capture relational properties. For example, formula

$$(\llbracket x \rrbracket_l < \llbracket x \rrbracket_s) \rightarrow \llbracket (\alpha, \beta) \rrbracket (\llbracket x \rrbracket_l < \llbracket x \rrbracket_s)$$

specifies that any pair of executions of α and β preserves the natural order on the values of x . (the projections $\llbracket \cdot \rrbracket_l$ and $\llbracket \cdot \rrbracket_s$ of variable x respectively refer to its value in the left and right executions).

We design the REL extension in a general and abstract manner, by focusing on the core constructs, i.e., programs and formulas, of all dynamic logics. It can be instantiated for different host dynamic logics to support domain-specific program constructs, i.e., constructs specific to the host logic. We conduct a case study by instantiating the REL extension for differential dynamic logic (dL) [34, 36], a well-known logic for verifying safety of cyber-physical systems. The instantiation is able to express key relational properties in a succinct way. Another case study on linear dynamic logic on finite traces [20] (LDL_f), a logic used by the AI community for reasoning about temporal constraints, can be found in the extended version of this paper.

Verifying REL formulas specified by REL instantiations can benefit from existing verification tools. We introduce an encoding of REL formulas that permit reusing existing tools to verify these formulas, inspired by the technique of self-composition [8], which reduces relational verification of a program to standard

³ *Color scheme:* we use blue color for syntactic constructs of host logics, and red color for those of the REL extension.

verification of a composition of two copies of the program. We develop a sound and complete encoding for the instantiation for $\mathbf{dL}$.

The REL extension naturally combines the general program constructs and domain-specific constructs, thereby promoting the verification of relational properties by integrating techniques for both general relational reasoning and domain-specific relational reasoning, i.e., relational reasoning on host logic specific constructs, e.g., $\mathbf{dL}$ programs modeling system dynamics. To investigate the usefulness of REL extension in promoting verification, we conduct an experiment on the instantiation for $\mathbf{dL}$. We implement several variants of its encoding and demonstrate that these variants can integrate techniques for general relational reasoning as well as domain-specific relational reasoning. We test these variants on a set of benchmarks and successfully verify all benchmarks using a theorem prover developed for $\mathbf{dL}$, with little to no manual effort. Some of the benchmarks cannot be automatically verified by prior tools to our knowledge.

Contribution. The key contribution of this paper is the design and validation of a general extension for dynamic logics to support lightweight yet effective first-class relational reasoning. In particular, we make the following contributions:

- A lightweight extension for dynamic logics to support first-class relational reasoning. Using the modalities, the REL extension can naturally and concisely express relational properties involving mixed quantifications. We present a set of practical example properties (Section 3).
- A case study that demonstrates the applicability of the REL extension. The instantiation can express key domain-specific relational properties that are difficult to express with the host logic (Section 4).
- An incrementally improvable encoding of REL formulas that can integrate general relational reasoning and domain-specific reasoning. We experiment with the encoding for $\mathbf{dL}$ on a set of benchmarks. The results are promising: we can use an existing tool for $\mathbf{dL}$ to verify non-trivial relational properties, with little to no manual effort. (Section 5).

Section 2 introduces the syntax and semantics of PDL. Section 6 discusses related work and Section 7 concludes.

2 Propositional Dynamic Logic

Propositional dynamic logic (PDL) is a subsystem of most, if not all, dynamic logics. It extends propositional logic with modalities to reason about program executions. The language of regular PDL has expressions of two sorts: (1) propositions or formulas and (2) programs. Its syntax is defined upon Π_0 , a set of atomic programs, and Φ_0 , a set of atomic propositions. Programs and propositions are mutually inductively defined from the atomic ones as shown in Figure 1. Programs include the operations of Kleene algebra with tests [30]: sequential composition, nondeterministic choice, nondeterministic repetition, and test of a formula. Formulas include the standard propositional connectives and program necessity $[\alpha]\phi$. Program existence $\langle\alpha\rangle\phi$ can be encoded with program necessity

and vice versa, e.g., $[\alpha]\phi = \neg\langle\alpha\rangle\neg\phi$. Common abbreviations for logical connectives apply.

The semantics of PDL formulas and programs is interpreted over a Kripke structure $(\mathcal{W}, \mathcal{R}, \mathcal{V})$, where $\mathcal{W}$ is a nonempty set of states, $\mathcal{R}$ is a mapping from the set Π_0 of atomic programs into binary relations on $\mathcal{W}$, and $\mathcal{V}$ is a mapping from the set Φ_0 of atomic propositions into subsets of $\mathcal{W}$. That is, $\mathcal{R}(P) \subseteq \mathcal{W} \times \mathcal{W}$ for $P \in \Pi_0$, and $\mathcal{V}(A) \subseteq \mathcal{W}$ for $A \in \Phi_0$. $\mathcal{R}$ and $\mathcal{V}$ are extended inductively to give meanings to all programs and formulas of PDL

as shown in Figure 1. We write $\omega \models \phi$ if formula ϕ is true at state ω , i.e., $\omega \in \mathcal{V}(\phi)$. We write $\llbracket \alpha \rrbracket$ to denote the semantics of α , i.e., if $(\omega, \nu) \in \llbracket \alpha \rrbracket$, then there is an execution of α that starts in state ω and ends in state ν .

Relations and states in PDL are *abstract*, i.e., states in PDL are abstract points and atomic programs in PDL are abstract binary relations. Such a level of abstraction lets us focus on the fundamental design of the REL extension.

3 The REL Extension

This section introduces the design of the REL extension in the setting of PDL. We present its syntax and semantics and then showcase example REL formulas.

3.1 Syntax and Semantics

The REL extension extends PDL with programs, modalities, and formulas, as follows. It builds on $\Pi_{\mathbf{r}}$, a set of REL atomic programs, and $\Phi_{\mathbf{r}}$, a set of REL atomic propositions. REL programs are analogous to PDL, with the addition of a biprogram construct (α, β) , where α and β specify two PDL programs to run, respectively, by the *left* and *right* execution. The syntax of REL formulas is also analogous to PDL, with the addition of projection formulas $[\phi]_{\mathbf{s}}$, which refers to a PDL formula ϕ in one of the two executions specified by $\mathbf{B} \in \{\mathbf{L}, \mathbf{R}\}$. The REL modality of necessity $\llbracket \alpha_{\mathbf{r}} \rrbracket \phi_{\mathbf{r}}$ can be used to encode the REL modality of existence $\langle \alpha_{\mathbf{r}} \rangle \phi_{\mathbf{r}}$ and vice versa, e.g., $\llbracket \alpha_{\mathbf{r}} \rrbracket \phi_{\mathbf{r}} = \neg \langle \alpha_{\mathbf{r}} \rangle \neg \phi_{\mathbf{r}}$.

The semantics of REL programs and formulas is interpreted over the semantics of REL atomic programs $\mathcal{R}_{\mathbf{r}}(P_{\mathbf{r}})$ (for $P_{\mathbf{r}} \in \Pi_{\mathbf{r}}$) and REL atomic propositions $\mathcal{V}_{\mathbf{r}}(A_{\mathbf{r}})$ (for $A_{\mathbf{r}} \in \Phi_{\mathbf{r}}$), as well as the semantics of PDL programs and formulas.

Program: $\alpha, \beta ::= P \in \Pi_0 \mid ?\phi \mid \alpha; \beta \mid \alpha \cup \beta \mid \alpha^*$

Formula: $\phi, \psi ::= \top \mid A \in \Phi_0 \mid \neg\phi \mid \phi \wedge \psi \mid [\alpha]\phi$

Semantics of formulas ϕ, ψ

$\omega \models \top$ iff $\omega \in \mathcal{W}$

$\omega \models A$ iff $\omega \in \mathcal{V}(A)$

$\omega \models \neg\phi$ iff $\omega \not\models \phi$

$\omega \models \phi \wedge \psi$ iff $\omega \models \phi$ and $\omega \models \psi$

$\omega \models [\alpha]\phi$ iff $\nu \models \phi$ for all states ν with $(\omega, \nu) \in \llbracket \alpha \rrbracket$

Semantics of programs α, β

$\llbracket P \rrbracket = \mathcal{R}(P)$ for $P \in \Pi_0$

$\llbracket ?\phi \rrbracket = \{(\omega, \omega) \mid \omega \models \phi\}$

$\llbracket \alpha; \beta \rrbracket = \{(\omega, \nu) \mid \exists \mu, (\omega, \mu) \in \llbracket \alpha \rrbracket \text{ and } (\mu, \nu) \in \llbracket \beta \rrbracket\}$

$\llbracket \alpha \cup \beta \rrbracket = \llbracket \alpha \rrbracket \cup \llbracket \beta \rrbracket$

$\llbracket \alpha^* \rrbracket = \llbracket \alpha \rrbracket^*$ the transitive, reflexive closure of $\llbracket \alpha \rrbracket$

Fig. 1: Syntax and Semantics of PDL

$\mathcal{R}_\tau(P_\tau)$ is a transition relation between two pairs of states: if $((\omega_l, \omega_r), (\nu_l, \nu_r)) \in \mathcal{R}_\tau(P_\tau)$, then an execution of the left program specified by P_τ runs from ω_l to ν_l , and an execution of the right program of P_τ runs from ω_r to ν_r . $\mathcal{V}_\tau(A_\tau)$ is a relation on states, i.e., $\mathcal{V}_\tau(A_\tau) \subseteq \mathcal{W} \times \mathcal{W}$ for $A_\tau \in \Phi_\tau$. Semantics of all REL programs and formulas, as shown in Figure 2, is defined inductively from the semantics of REL atomic programs and formulas, and the projection formulas. A projection formula $\lfloor \phi \rfloor_\mathbf{B}$ tests if PDL formula ϕ holds in the state specified by $\mathbf{B}$ (left or right). REL modality of necessity $\llbracket \alpha_\tau \rrbracket \phi_\tau$ holds in bi-state (ω_l, ω_r) if and only if formula ϕ_τ holds in any pairs of states that are reachable from (ω_l, ω_r) by running α_τ . We write $\llbracket \alpha_\tau \rrbracket_{\text{RL}}$ to denote the semantics of α_τ , and $(\omega_l, \omega_r) \models_{\text{RL}} \phi_\tau$ if the formula ϕ_τ holds in bi-state (ω_l, ω_r) .

$$\begin{aligned} \text{REL program: } \alpha_\tau, \beta_\tau &::= P_\tau \in \Pi_\tau \mid ?\phi_\tau \mid \alpha_\tau; \beta_\tau \mid \alpha_\tau \cup \beta_\tau \mid \alpha_\tau^* \mid (\alpha, \beta) \\ \text{REL formula: } \phi_\tau, \psi_\tau &::= A_\tau \mid \neg\phi_\tau \mid \phi_\tau \wedge \psi_\tau \mid \llbracket \alpha_\tau \rrbracket \phi_\tau \mid \lfloor \phi \rfloor_\mathbf{B} \quad (\mathbf{B} \in \{\mathbf{L}, \mathbf{R}\}) \end{aligned}$$

Semantics of REL program $\alpha_\tau \beta_\tau$

$$\begin{aligned} \llbracket P_\tau \rrbracket_{\text{RL}} &= \mathcal{R}_\tau(P_\tau) \text{ for } P_\tau \in \Pi_\tau \\ \llbracket ?\phi_\tau \rrbracket_{\text{RL}} &= \{((\omega_l, \omega_r), (\omega_l, \omega_r)) \mid (\omega_l, \omega_r) \models_{\text{RL}} \phi_\tau\} \\ \llbracket \alpha_\tau; \beta_\tau \rrbracket_{\text{RL}} &= \{((\omega_l, \omega_r), (\nu_l, \nu_r)) \mid \text{exists } \mu_l, \mu_r \text{ that} \\ &\quad ((\omega_l, \omega_r), (\mu_l, \mu_r)) \in \llbracket \alpha_\tau \rrbracket_{\text{RL}} \text{ and } ((\mu_l, \mu_r), (\nu_l, \nu_r)) \in \llbracket \beta_\tau \rrbracket_{\text{RL}}\} \\ \llbracket \alpha_\tau \cup \beta_\tau \rrbracket_{\text{RL}} &= \{((\omega_l, \omega_r), (\nu_l, \nu_r)) \mid \llbracket \alpha_\tau \rrbracket_{\text{RL}} \cup \llbracket \beta_\tau \rrbracket_{\text{RL}}\} \\ \llbracket \alpha_\tau^* \rrbracket_{\text{RL}} &= \llbracket \alpha_\tau \rrbracket_{\text{RL}}^* \text{ the transitive, reflexive closure of } \llbracket \alpha_\tau \rrbracket_{\text{RL}}, \text{ where } \alpha_\tau^0 \text{ is } (?T, ?T) \\ \llbracket (\alpha, \beta) \rrbracket_{\text{RL}} &= \{((\omega_l, \omega_r), (\nu_l, \nu_r)) \mid (\omega_l, \nu_l) \in \llbracket \alpha \rrbracket \text{ and } (\omega_r, \nu_r) \in \llbracket \beta \rrbracket\} \end{aligned}$$

Semantics of REL formula $\phi_\tau \psi_\tau$

Truth of formula ϕ_τ in bi-state ω_l, ω_r , denoted $(\omega_l, \omega_r) \models_{\text{RL}} \phi_\tau$ is defined as follows:

$$\begin{aligned} (\omega_l, \omega_r) \models_{\text{RL}} A_\tau &\text{ iff } (\omega_l, \omega_r) \in \mathcal{V}_\tau(A_\tau) \\ (\omega_l, \omega_r) \models_{\text{RL}} \neg\phi_\tau &\text{ iff } (\omega_l, \omega_r) \not\models_{\text{RL}} \phi_\tau \\ (\omega_l, \omega_r) \models_{\text{RL}} \phi_\tau \wedge \psi_\tau &\text{ iff } (\omega_l, \omega_r) \models_{\text{RL}} \phi_\tau \text{ and } (\omega_l, \omega_r) \models_{\text{RL}} \psi_\tau \\ (\omega_l, \omega_r) \models_{\text{RL}} \llbracket \alpha_\tau \rrbracket \phi_\tau &\text{ iff } \forall \nu_l, \nu_r \text{ that } ((\omega_l, \omega_r), (\nu_l, \nu_r)) \in \llbracket \alpha_\tau \rrbracket_{\text{RL}}, (\nu_l, \nu_r) \models_{\text{RL}} \phi_\tau \text{ holds} \\ (\omega_l, \omega_r) \models_{\text{RL}} \lfloor \phi \rfloor_\mathbf{B} &\text{ iff } \omega_l \models \phi \text{ when } \mathbf{B} = \mathbf{L} \text{ or } \omega_r \models \phi \text{ when } \mathbf{B} = \mathbf{R} \end{aligned}$$

Fig. 2: Syntax and semantics of REL programs and formulas

3.2 Quantification over Executions

Expressing relational properties often involves quantifications over executions, e.g., for all executions of a program α , there exists an execution of β that gets the same results. We can use REL formulas to easily and *succinctly* express all quantifications over two executions by using (and mixing) the necessity and existence modalities, as shown below. We write σ_α (or σ_β) to denote an execution of program α (or β) and $\sigma_\alpha[-1]$ to denote the last state of σ_α . The first two entries show how to encode quantifications without alternations. The other four show how to encode quantifications with alternations, by combining the biprogram construct and a special program $?T$ (i.e., *skip* in common programming languages). For example, *forall, exists*, i.e., $\forall \sigma_\alpha \exists \sigma_\beta$ is encoded as $\llbracket (\alpha, ?T) \rrbracket \llbracket \langle ?T, \beta \rangle \rrbracket$, where the modality $\llbracket (\alpha, ?T) \rrbracket$ quantifies over executions of program α , and the modality $\llbracket \langle ?T, \beta \rangle \rrbracket$ quantifies over executions of program β .

$$\begin{aligned}
\forall \sigma_\alpha \forall \sigma_\beta. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \llbracket (\alpha, \beta) \rrbracket \phi_\tau \\
\exists \sigma_\alpha \exists \sigma_\beta. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \langle\langle \alpha, \beta \rangle\rangle \phi_\tau \\
\forall \sigma_\alpha \exists \sigma_\beta. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \llbracket (\alpha, ?\top) \rrbracket \langle\langle ?\top, \beta \rangle\rangle \phi_\tau \\
\forall \sigma_\beta \exists \sigma_\alpha. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \llbracket (? \top, \beta) \rrbracket \langle\langle \alpha, ?\top \rangle\rangle \phi_\tau \\
\exists \sigma_\alpha \forall \sigma_\beta. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \langle\langle \alpha, ?\top \rangle\rangle \llbracket (? \top, \beta) \rrbracket \phi_\tau \\
\exists \sigma_\beta \forall \sigma_\alpha. ((\sigma_\alpha[-1], \sigma_\beta[-1]) \models_{\text{RL}} \phi_\tau) &\equiv \langle\langle ?\top, \beta \rangle\rangle \llbracket (\alpha, ?\top) \rrbracket \phi_\tau
\end{aligned}$$

Fig. 3: Encoding of quantifications

Notation: For brevity, we write $\llbracket (\alpha, \beta) \rrbracket \phi_\tau$ and $\langle\langle \alpha, \beta \rangle\rangle \phi_\tau$ for formulas $\llbracket (\alpha, ?\top) \rrbracket \langle\langle ?\top, \beta \rangle\rangle \phi_\tau$, and $\langle\langle \alpha, ?\top \rangle\rangle \llbracket (? \top, \beta) \rrbracket \phi_\tau$ respectively from now on. We can design other syntactic sugar to further improve the usability of REL.

3.3 Expressing Relational Properties

The REL extensions express relational properties intuitively. We show some interesting examples here. Some examples involve syntax more than PDL, but their meanings can be easily inferred. Note that example properties shown here are general: they can be similarly expressed in most dynamic logics.

Refinement Relation. A program α refines another program β if the behaviors of β subsume those of α . That is, all states reachable from a state ω by following a transition of α could also be reached from ω by following *some* transitions of β . Such a property can be encoded using the $\forall \sigma_\alpha, \exists \sigma_\beta$ quantification:

$$(\llbracket A \rrbracket_l \leftrightarrow \llbracket A \rrbracket_r) \rightarrow \llbracket (\alpha, \beta) \rrbracket (\llbracket A \rrbracket_l \leftrightarrow \llbracket A \rrbracket_r)$$

Where programs α and β refer to the same set of atomic propositions and programs. The formula $(\llbracket A \rrbracket_l \leftrightarrow \llbracket A \rrbracket_r)$ is a shorthand for $\bigwedge_{i \in 1 \dots n} (\llbracket A_i \rrbracket_l \leftrightarrow \llbracket A_i \rrbracket_r)$, which encodes that the left and right states are identical.

Noninterference. Noninterference [39, 25] is a well-known strong information security property that guarantees that public outputs of a system do not reveal any confidential information (i.e., confidentiality), or dually that untrusted inputs of a system do not modify trusted contents (i.e., integrity). Noninterference has various variants, especially in a language that involves nondeterminism. Here, we express a few variants.

First, a common notion of noninterference for confidentiality of a *deterministic* program ensures that attackers who have access to the program's low-security input and output won't be able to infer the program's high-security input. Intuitively, the property states that for any executions of a deterministic program α receiving the same low-security inputs, they should produce the same low-security outputs. Such a noninterference notion is a $\forall \forall$ property:

$$(\llbracket A_{in} \rrbracket_l \leftrightarrow \llbracket A_{in} \rrbracket_r) \rightarrow \llbracket (\alpha, \alpha) \rrbracket (\llbracket A_{out} \rrbracket_l \leftrightarrow \llbracket A_{out} \rrbracket_r)$$

Where atomic propositions A_{in} and A_{out} represent, respectively, low-security inputs and outputs of α . The modality $\llbracket (\alpha, \alpha) \rrbracket$ ensures that α produces the same low-security output if the same low-security input is received.

Various forms of noninterference can be defined when α has nondeterminism. One of them is *possibilistic noninterference* [39] or *nondeducibility* [2], which states that an attacker cannot infer confidential information *for certain* in a program with non-observable nondeterminism. Intuitively, the attacker cannot

observe the nondeterministic choices made by program α , so the attacker is not certain if the information received is confidential. Such a noninterference notion can be expressed with the *forall*, *exists* quantifications over executions of α , i.e.,

$$(\llbracket A_{in} \rrbracket_l \leftrightarrow \llbracket A_{in} \rrbracket_r) \rightarrow \llbracket (\alpha, \alpha) \rrbracket (\llbracket A_{out} \rrbracket_l \leftrightarrow \llbracket A_{out} \rrbracket_r)$$

To properly reason about two executions in the presence of nondeterminism, some variants of noninterference require (some of) the nondeterminism in the two executions to resolve in the same way [48]. For example, a nondeterministic choice in a program may represent a user's decision, which is assumed to be public input, and so the resolution of the nondeterministic choice should be the same in both executions. Under such circumstances, we can use the REL nondeterministic choice, i.e., $\alpha_r \cup \beta_r$, to express that two executions take the same choice. The following formula shows an example of such a noninterference, where α_r is $((?A_{in}; ?A_{out}, ?A_{in}; ?A_{out}))$ and β_r is $((?\neg A_{out}, ?\neg A_{out}))$:

$$(\llbracket A_{in} \rrbracket_l \leftrightarrow \llbracket A_{in} \rrbracket_r) \rightarrow \llbracket \alpha_r \cup \beta_r \rrbracket (\llbracket A_{out} \rrbracket_l \leftrightarrow \llbracket A_{out} \rrbracket_r)$$

This formula holds: two executions with the same nondeterministic choice, i.e., either the first branch or the second, would lead to the same value of A_{out} .

Lock-step Properties. Note that the semantics of α_r^* means the left and right programs specified in α_r would always loop for the same number of iterations. Combining it with other constructs, we can express interesting relational properties on loops that involve lock-step. For example, the following formula expresses a *relational invariant* of a loop with a biprogram. In particular, an invariant ϕ_{inv} holds at the end of every loop iteration of $(\alpha, \beta)^*$: $\phi_{inv} \rightarrow \llbracket (\alpha, \beta) \rrbracket \phi_{inv}$.

Loop Alignments. As shown in prior work, verifying relational properties on loops may require careful *alignment* in executing the left and right loops. The REL extension can readily and clearly express these alignments. Consider the example (DoubleSquare) adapted from [41, 3] on integers x , y , and z :

$$\begin{aligned} \llbracket x \rrbracket_l = \llbracket x \rrbracket_r \rightarrow \llbracket y := 0; (z := 2x, z := x); (body^*, body^*) \rrbracket \\ \llbracket (?z \leq 0, ?z \leq 0); (? \top, y := 2y) \rrbracket \llbracket y \rrbracket_l = \llbracket y \rrbracket_r \end{aligned}$$

Where $body \equiv ?z \geq 0; z := z - 1; y := y + x$. An effective alignment for verifying this property can be expressed as a formula that replaces $(body^*, body^*)$ with $(body; body, body)^*$ which aligns two loop iterations of the left execution with one iteration of the right.

For general $\forall\forall$ formulas of the form $\llbracket (body^*, body^*) \rrbracket \phi_r$, we can express the alignments required for verifying such a formula in a general manner, as follows:

$$\llbracket (?Gb; (body, body) \cup ?Gl; (body, \top) \cup ?Gr; (\top, body))^* \rrbracket \phi_r$$

Where formulas Gb , Gl , and Gr specify the conditions for executing, respectively, both the left and right loops, the left loop only, and the right loop only. For example, we can get an effective alignment for verifying the example above by letting $Gb \equiv \llbracket x \rrbracket_l = \llbracket x \rrbracket_r$, $Gl \equiv \llbracket x \rrbracket_l \geq \llbracket x \rrbracket_r$, $Gr \equiv \llbracket false \rrbracket_l \wedge \llbracket false \rrbracket_r$.

Motion Planning. The REL extension can intuitively express many relational properties on motion planning or path finding, e.g., for driving or robots. For example, an important property in motion planning is finding the optimal path (of all possible paths). Given a control system modeled as α , the following $\exists\forall$

formula specifies that an optimal path exists (the left execution arrives no-later than all possible executions):

$$([\phi]_l \wedge [\phi]_r) \rightarrow \langle (\alpha, \alpha) \rangle ([\psi]_r \rightarrow [\psi]_l)$$

Where $[\phi]_l \wedge [\phi]_r$ specifies that both executions start at the same location, and $[\psi]_r \rightarrow [\psi]_l$ specifies that the left execution arrives no-later than other executions of α . We can further express useful variants of this property. For example, an optimal path may exist only if certain actions are triggered:

$$([\phi]_l \wedge [\phi]_r) \rightarrow \langle ((\alpha_1; \beta; \alpha_2), \alpha) \rangle ([\psi]_l \rightarrow [\psi]_r)$$

Where β specifies these actions. And α_1 and α_2 respectively model possible actions the system can take before and after β .

4 Case Study: Differential Dynamic Logic

This section presents a case study of the REL extension on differential dynamic logic [34, 37, 36] ($\text{d}\mathcal{L}$), a well-known first-order dynamic logic that enables verifying high-level cyber-physical systems models featuring real arithmetic, non-determinism, and differential equations. Programs in $\text{d}\mathcal{L}$ are referred to as *hybrid programs* [37]. They are a formalism for modeling systems that have both continuous and discrete behaviors. Hybrid programs can express continuous evolution (as differential equations) as well as discrete transitions.

Figure 4 shows the syntax of constructs that $\text{d}\mathcal{L}$ add over PDL and their semantics. Variables are real-valued and can be deterministically assigned ($x := \theta$, where θ is an arithmetic term) or nondeterministically assigned ($x := *$). The hybrid program $x' = \theta \& \phi$ expresses the continuous evolution of variables x : given the current value of variable x , the system follows the differential equation $x' = \theta$ for some (nondeterministically chosen) amount of time so long as the formula ϕ , the *evolution domain constraint*, holds for all of that time. Note that x can be a vector of variables and then θ is a vector of terms of the same dimension. Atomic formulas of $\text{d}\mathcal{L}$ are comparisons of terms, i.e., $(\theta \sim \delta)$.

Similar to PDL, the semantics of $\text{d}\mathcal{L}$ [34, 36] is a Kripke semantics where the Kripke model's worlds are the states of the system. Let $\mathbb{R}$ denote the set of real numbers and $\mathbb{V}$ denote the set of variables. A state is a map $\omega : \mathbb{V} \mapsto \mathbb{R}$ assigning a real value $\omega(x)$ to each variable $x \in \mathbb{V}$. The set of all states is denoted by STA.

With $\text{d}\mathcal{L}$, we are often interested in formulas of the form $\phi_{pre} \rightarrow [\alpha] \phi_{post}$: if ϕ_{pre} is true then ϕ_{post} holds after any possible execution of α . The hybrid program α often has the form $(ctrl; plant)^*$, where *ctrl* models atomic actions of the control system and does not contain continuous parts (i.e., differential equations); and *plant* models the evolution of the physical environment and has the form of $x' = \theta \& \phi$. That is, the system is modeled as unbounded repetitions of a controller action followed by an update to the physical environment.

Consider, as an example, an autonomous vehicle that needs to stop before hitting an obstacle [37]. Figure 5 shows a $\text{d}\mathcal{L}$ model of such an autonomous vehicle. Let d be the vehicle's distance from the obstacle. The desired *safety condition* (ϕ_{post}) is that d is positive, i.e., not hitting the obstacle. Let v and

Semantics of dL terms $\theta, \delta ::= x \mid c \mid \theta \oplus \delta$

$$\omega \llbracket x \rrbracket = \omega(x) \text{ and } \omega \llbracket c \rrbracket = c$$

$$\omega \llbracket \theta \oplus \delta \rrbracket = \omega \llbracket \theta \rrbracket \oplus \omega \llbracket \delta \rrbracket \text{ (} \oplus \text{ denotes arithmetic for } \oplus \in \{+, \times\} \text{)}$$

Semantics of dL programs $\alpha, \beta ::= x := \theta \mid x := * \mid x' = \theta \& \phi \mid \dots$

$$\llbracket x := \theta \rrbracket = \{(\omega, \nu) \mid \nu(x) = \omega \llbracket \theta \rrbracket \text{ and for all other variables } y \neq x, \nu(y) = \omega(y)\}$$

$$\llbracket x := * \rrbracket = \{(\omega, \nu) \mid \nu(y) = \omega(y) \text{ for all variables } y \neq x\}$$

$$\llbracket x' = \theta \& \phi \rrbracket = \{(\omega, \nu) \mid \text{exists a solution } \varphi : [0, r] \rightarrow \text{STA of } x' = \theta \text{ with } \varphi(0) = \omega, \varphi(r) = \nu, \text{ and } \varphi(t) \models \phi \text{ for all } t \in [0, r]\}$$

Semantics of dL formulas $\phi, \psi ::= \theta \sim \delta \mid \forall x. \phi \mid \dots$

$$\omega \models \theta \sim \delta \text{ iff } \omega \llbracket \theta \rrbracket \sim \omega \llbracket \delta \rrbracket \text{ (} \sim \text{ denotes comparison for } \sim \in \{=, \leq, <, \geq, >\}$$

$$\omega \models \forall x. \phi \text{ iff } \nu \models \phi \text{ for all states } \nu \text{ that agree with } \omega \text{ except for the value of } x$$

Fig. 4: Syntax and semantics of dL programs and formulas in addition to PDL

a be the vehicle's velocity and acceleration. Let t be the time elapsed since the controller was last invoked. The program *plant* describes how the physical environment evolves over time interval ϵ : distance changes according to $-v$ (i.e., $d' = -v$), velocity changes according to the acceleration (i.e., $v' = a$), and time passes at a constant rate (i.e., $t' = 1$). The differential equations evolve within the time interval $t \leq \epsilon$ and if v is non-negative (i.e., $v \geq 0$).

$$\begin{aligned} \phi_{pre} &\equiv A \geq 0 \wedge B \geq 0 \wedge 2Bd > v^2 \\ \phi_{post} &\equiv d > 0 \\ \psi &\equiv 2Bd > v^2 + (A + B)(A\epsilon^2 + 2v\epsilon) \\ accel &\equiv ?\psi; a := A \\ brake &\equiv a := -B \\ ctrl &\equiv (accel \cup brake); t := 0 \\ plant &\equiv d' = -v, v' = a, t' = 1 \& (v \geq 0 \wedge t \leq \epsilon) \\ \phi_{safety} &\equiv \phi_{pre} \rightarrow [(ctrl; plant)^*] \phi_{post} \end{aligned}$$

Fig. 5: dL model of an autonomous vehicle

The hybrid program *ctrl* models the vehicle's controller. The vehicle can either accelerate at A or brake at $-B$. The controller chooses nondeterministically between these options. Programs *accel* and *brake* express the controller accelerating or braking. The controller can accelerate only if ψ is true, which captures that the vehicle can accelerate for the

next ϵ seconds only if doing so would still allow it to brake in time. The formula to be verified, ϕ_{safety} , is shown at the last line of Figure 5. Given a precondition ϕ_{pre} , the axioms and proof rules of dL can be used to prove the safety condition ϕ_{post} . The theorem prover KeYmaera X [23] provides tool support.

4.1 dL_{REL}: A REL Instantiation for dL

We build dL_{REL}, a REL instantiation for dL. dL_{REL} instantiates the atomic programs and formulas of the REL extension with constructs designed for dL. The syntax and semantics of these constructs are shown in Figure 6. (We omit the other constructs shown in Figure 2.) The relational atomic programs, i.e., P_r , in dL_{REL} are a relational version of deterministic assignment ($x := \theta$), nondeterministic assignment ($x := *$), and continuous evolution ($x' = \theta \& \phi$). The relational atomic formulas, i.e., A_r , in dL_{REL} are comparisons of dL_{REL} terms, i.e., $\theta_r \sim \delta_r$. A dL_{REL} term can be either a *projection* term: $\llbracket \theta \rrbracket_s$, which refers to a dL term θ in one of the two executions specified by $B \in \{L, R\}$, or an arithmetic operation of two dL_{REL} terms, i.e., $\theta_r \oplus \delta_r$.

Semantics of $\mathbf{dL}_{\text{REL}}$ term: $\theta_{\tau}, \delta_{\tau} ::= \lfloor \theta \rfloor_{\mathbf{B}} \mid \theta_{\tau} \oplus \delta_{\tau}$
 $(\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \lfloor \theta \rfloor_{\mathbf{B}} \rrbracket_{\text{RL}} = \omega_{\mathbf{L}} \llbracket \theta \rrbracket$ if $\mathbf{B} = \mathbf{L}$ or $\omega_{\mathbf{R}} \llbracket \theta \rrbracket$ if $\mathbf{B} = \mathbf{R}$
 $(\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \theta_{\tau} \oplus \delta_{\tau} \rrbracket_{\text{RL}} = (\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \theta_{\tau} \rrbracket_{\text{RL}} \oplus (\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \delta_{\tau} \rrbracket_{\text{RL}}$ ($\oplus$ denotes arithmetic for $\oplus \in \{+, \times\}$)
Semantics of $\mathbf{dL}_{\text{REL}}$ program: $\alpha_{\tau}, \beta_{\tau} ::= x := \theta \mid x := * \mid x' = \theta \& \phi \mid \dots$
 $\llbracket x := \theta \rrbracket_{\text{RL}} = \llbracket (x := \theta, x := \theta) \rrbracket_{\text{RL}}$
 $\llbracket x := * \rrbracket_{\text{RL}} = \llbracket (x := *, x := *) ; ?(\lfloor x \rfloor_{\mathbf{L}} = \lfloor x \rfloor_{\mathbf{R}}) \rrbracket_{\text{RL}}$
 $\llbracket x' = \theta \& \phi \rrbracket_{\text{RL}} = \llbracket (t := 0, t := 0) ; (x' = \theta, t' = 1 \& \phi, x' = \theta, t' = 1 \& \phi) ; ?(\lfloor t \rfloor_{\mathbf{L}} = \lfloor t \rfloor_{\mathbf{R}}) \rrbracket_{\text{RL}}$
Semantics of $\mathbf{dL}_{\text{REL}}$ formula: $\phi_{\tau}, \psi_{\tau} ::= \theta_{\tau} \sim \delta_{\tau} \mid \dots$
 $(\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \models_{\text{RL}} \theta_{\tau} \sim \delta_{\tau}$ iff $(\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \theta_{\tau} \rrbracket_{\text{RL}} \sim (\omega_{\mathbf{L}}, \omega_{\mathbf{R}}) \llbracket \delta_{\tau} \rrbracket_{\text{RL}}$ ($\sim$ is arithmetic comparison of $\sim$)
 Fig. 6: Syntax and semantics of $\mathbf{dL}_{\text{REL}}$ programs and formulas in addition to REL

Semantically, program $\llbracket x := \theta \rrbracket_{\text{RL}}$ indicates that the same deterministic assignment is run by the left and right executions. The program $\llbracket x := * \rrbracket_{\text{RL}}$ ensures both executions having the same value for variable x . $\llbracket x' = \theta \& \phi \rrbracket_{\text{RL}}$ enforces a constraint that the *durations* used by the physical evolution in both executions are the same (i.e., $?(\lfloor t \rfloor_{\mathbf{L}} = \lfloor t \rfloor_{\mathbf{R}})$ where t is a fresh continuous variable added into both executions). Such a design is useful since we often want to compare executions of two systems only if they execute for the same period of duration.

$\psi \equiv 2Bd > v^2 + (A + B)(A\epsilon^2 + 2v\epsilon)$
 $\text{accel} \equiv (? \psi, ? \psi) ; a := A$
 $\text{brake} \equiv a := -B$
 $\text{ctrl}_t \equiv (\quad (? \text{temp} > T ; \text{thermo} := -1) \quad \cup \quad (? \text{temp} < T ; \text{thermo} := 1) \quad \cup \quad (? \text{temp} = T) \quad)$
 $\text{ctrl}'_t \equiv (\quad (? \text{temp} > T ; \text{thermo} := -2) \quad \cup \quad (? \text{temp} < T ; \text{thermo} := 2) \quad \cup \quad (? \text{temp} = T) \quad)$
 $\text{ctrl}_v \equiv (\text{accel} \cup \text{brake}) ; t := 0$
 $\text{plant} \equiv d' = -v, v' = a, t' = 1, \text{temp}' = \text{thermo} \quad \& \quad (v \geq 0 \wedge t \leq \epsilon)$
 $\phi_{\tau} \equiv (\lfloor d \rfloor_{\mathbf{L}} = \lfloor d \rfloor_{\mathbf{R}} \wedge \lfloor v \rfloor_{\mathbf{L}} = \lfloor v \rfloor_{\mathbf{R}})$
 $\phi_{\text{robust}} \equiv \phi_{\tau} \rightarrow \llbracket ((\text{ctrl}_t, \text{ctrl}'_t) ; \text{ctrl}_v ; \text{plant})^* \rrbracket \phi_{\tau}$

Fig. 7: A $\mathbf{dL}_{\text{REL}}$ example on robustness

temperature T , and then set the thermostat accordingly. The modes differ only in the values of thermo . The temperature changes according to thermo .

A system designer may want a *robust* velocity control, i.e., the choice of modes for temperature control won't interfere with the vehicle's control of velocity. This relational property is expressed as a $\mathbf{dL}_{\text{REL}}$ formula at the last line (i.e., ϕ_{robust}). Intuitively, the formula says for two runs of the vehicle, that have different modes of temperature control, if the vehicle starts with the same position and velocity (the premise ϕ_{τ} of the implication), makes the same control decisions for acceleration and brake (ctrl_v), and runs for the same duration (plant), it would end with the same position and velocity (the conclusion ϕ_{τ} of the implication). The validity of this formula suggests that the vehicle has robust velocity control.

Expressing such a relational property with $\mathbf{dL}_{\text{REL}}$ is straightforward and much more succinct than the original version [48].

Figure 7 shows examples of $\mathbf{dL}_{\text{REL}}$ programs and formulas adapted from existing work [48]. It presents a design of an autonomous vehicle with velocity control and interior temperature control. Its velocity control is the same as the example presented in Figure 5. For temperature control, the vehicle detects the interior temperature (temp), and then chooses one of the two control modes, specified respectively in program ctrl_t and ctrl'_t . Here, both modes compare the current temperature with a target temperature

We investigate the expressive power of $\mathbf{dL}_{\text{REL}}$ and prove that it is equally expressive as $\mathbf{dL}$, stated with the following theorem:

Theorem 1. *$\mathbf{dL}_{\text{REL}}$ and $\mathbf{dL}$ are equally expressive.*

The theorem holds as every $\mathbf{dL}$ formula can be encoded as an equivalent $\mathbf{dL}_{\text{REL}}$ formula and vice versa. We prove this equivalence by constructing a sound and complete encoding of $\mathbf{dL}_{\text{REL}}$ in $\mathbf{dL}$ in Section 5. Such a $\mathbf{dL}_{\text{REL}}$ design promotes verification, as all $\mathbf{dL}_{\text{REL}}$ formulas can be verified using the tools for $\mathbf{dL}$ and leverage their future advances.

Note that it's possible to make a different design so $\mathbf{dL}_{\text{REL}}$ is more expressive. For example, we could introduce new atomic relational program constructs that cannot be expressed with $\mathbf{dL}$. For example, recent work [14, 49] introduce a notion of simulation distance in the setting of $\mathbf{dL}$, a relational property that quantitatively captures the behavior distance between two $\mathbf{dL}$ programs. The notion requires calculating the *infimum* of distances between a reachable state of one program to the set of reachable states of another program. We can add a new atomic relational formula to help express such a relational property, and such a atomic construct cannot be expressed with $\mathbf{dL}$ due to the infimum calculation.

5 Verification Techniques

Verifying REL formulas benefits from existing tools. To explore this benefit, we introduce an encoding that transforms REL formulas into formulas of the host logic, so existing tools and techniques developed for the host logic can be reused. The encoding can integrate verification techniques for general relational reasoning as well as domain-specific reasoning. Thanks to the integration, it promotes automated verification of relational properties, especially domain-specific ones. To demonstrate its usefulness, we conduct an experiment with the encoding for $\mathbf{dL}_{\text{REL}}$ on a set of benchmarks. The results show that we can verify the benchmarks with an existing tool developed for $\mathbf{dL}$, with little to no manual effort.

5.1 Host Logic Encoding

The encoding is inspired by self-composition [22, 8, 45, 33], a proof technique often used for proving noninterference for deterministic programs. We explain self-composition first, and then introduce the encoding in the setting of $\mathbf{dL}_{\text{REL}}$.

Self-Composition. To develop an intuition for how the self-composition technique is used to prove noninterference, consider the problem of checking whether low-security outputs of a deterministic program reveal high-security inputs. Construct two copies of the program, renaming the program variables so that the variables in the two copies are disjoint. Set the low-security inputs in both copies to identical values but allow the high-security inputs to take different values. Now, sequentially compose these two programs together. If the composed program can terminate in a state where the corresponding low-security outputs

differ, then the original program does not satisfy noninterference; conversely, if in all executions of the composed program, the low-security outputs are the same, then the original program satisfies noninterference. Intuitively, the composition of the two copies allows a single program to represent two executions of the original program, reducing checking a relational property of the original problem to checking a safety property of the composed program.

Using the insights, our encoding for $\mathbf{dL}_{\text{REL}}$ builds on two steps: (1) renaming a $\mathbf{dL}_{\text{REL}}$ formula into an equivalent one whose left and right programs use disjoint variables, and (2) making a composition of the two programs.

Renaming $\mathbf{dL}_{\text{REL}}$ Formulas. To help with renaming, we define *renaming functions* that map all variables accessible by the *right* execution/program to fresh variables. We write $\xi(\phi_{\mathbf{r}})$ for the formula identical to $\phi_{\mathbf{r}}$ but whose variables accessible by the right execution have been renamed according to ξ . Renaming functions similarly apply to $\mathbf{dL}_{\text{REL}}$ programs. We also write $\xi(\omega)$ for the state identical to ω but whose domain variables have been renamed according to ξ .

A translation from $\mathbf{dL}_{\text{REL}}$ to $\mathbf{dL}$. We develop a function π that will be used to transform renamed $\mathbf{dL}_{\text{REL}}$ formulas, i.e., $\xi(\phi_{\mathbf{r}})$, into an equivalent $\mathbf{dL}$ formula. A key step of π is to convert the biprogram construct (α, β) into a composition of α and β . Also, it directly extracts the contents from all projection constructs, e.g., $\lfloor \phi \rfloor_{\mathbf{s}}$ and $\lfloor \theta \rfloor_{\mathbf{s}}$. The π is defined as follows:

Definition 1 (π for $\mathbf{dL}_{\text{REL}}$). For a $\mathbf{dL}_{\text{REL}}$ formula $\phi_{\mathbf{r}}$ whose variables for the left and right executions are disjoint (i.e., $\text{VAR}_l(\phi_{\mathbf{r}}) \cap \text{VAR}_r(\phi_{\mathbf{r}}) = \emptyset$), a function π that can transform a $\mathbf{dL}_{\text{REL}}$ formula $\phi_{\mathbf{r}}$ (or a $\mathbf{dL}_{\text{REL}}$ program $\alpha_{\mathbf{r}}$) to an equivalent $\mathbf{dL}$ formula (or program), is defined inductively as follows.

$$\begin{aligned} \pi(\theta_{\mathbf{r}} \sim \delta_{\mathbf{r}}) &= \pi(\theta_{\mathbf{r}}) \sim \pi(\delta_{\mathbf{r}}) & \pi(\alpha_{\mathbf{r}} ; \beta_{\mathbf{r}}) &= \pi(\alpha_{\mathbf{r}}) ; \pi(\beta_{\mathbf{r}}) \\ \pi(\neg \phi_{\mathbf{r}}) &= \neg \pi(\phi_{\mathbf{r}}) & \pi(\alpha_{\mathbf{r}} \cup \beta_{\mathbf{r}}) &= \pi(\alpha_{\mathbf{r}}) \cup \pi(\beta_{\mathbf{r}}) \\ \pi(\phi_{\mathbf{r}} \wedge \psi_{\mathbf{r}}) &= \pi(\phi_{\mathbf{r}}) \wedge \pi(\psi_{\mathbf{r}}) & \pi(? \phi_{\mathbf{r}}) &= ? \pi(\phi_{\mathbf{r}}) \\ \pi(\llbracket \alpha_{\mathbf{r}} \rrbracket \phi_{\mathbf{r}}) &= \llbracket \pi(\alpha_{\mathbf{r}}) \rrbracket \pi(\phi_{\mathbf{r}}) & \pi(\alpha_{\mathbf{r}}^*) &= (\pi(\alpha_{\mathbf{r}}))^* \\ \pi(\lfloor \phi \rfloor_{\mathbf{s}}) &= \phi & \pi((\alpha, \beta)) &= \alpha ; \beta \\ \text{And } \pi(\theta_{\mathbf{r}}) \text{ on terms is inductively defined:} & & \pi(\theta_{\mathbf{r}} \oplus \delta_{\mathbf{r}}) &= \pi(\theta_{\mathbf{r}}) \oplus \pi(\delta_{\mathbf{r}}) \\ \pi(\lfloor \theta \rfloor_{\mathbf{s}}) &= \theta & & \end{aligned}$$

With a renaming function ξ , we can encode a $\mathbf{dL}_{\text{REL}}$ formula $\phi_{\mathbf{r}}$ as $\pi(\xi(\phi_{\mathbf{r}}))$.

Theorem 2 (Soundness and completeness of the encoding for $\mathbf{dL}_{\text{REL}}$).

For states ω_l, ω_r , a $\mathbf{dL}_{\text{REL}}$ formula $\phi_{\mathbf{r}}$, and a renaming function ξ for $\phi_{\mathbf{r}}$,

$$(\omega_l, \omega_r) \models_{\text{REL}} \phi_{\mathbf{r}} \text{ iff } ((\omega_l \Downarrow \text{VAR}_l(\phi_{\mathbf{r}})) \otimes (\xi(\omega_r \Downarrow \text{VAR}_r(\phi_{\mathbf{r}})))) \models \pi(\xi(\phi_{\mathbf{r}}))$$

Where $\omega \Downarrow V$ denotes the projection of state ω on a set of variables V : the map $\{x \mapsto \omega(x)\}$ for all $x \in V$, and $\otimes$ means the join of two non-overlapping states.

The proof can be done by simultaneous induction on $\phi_{\mathbf{r}}$ and $\alpha_{\mathbf{r}}$.

Verifying $\mathbf{dL}_{\text{REL}}$ Formulas with existing tools. With Theorem 2, we can verify a $\mathbf{dL}_{\text{REL}}$ formula $\phi_{\mathbf{r}}$ by verifying a $\mathbf{dL}$ formula $\pi(\xi(\phi_{\mathbf{r}}))$, which can be done with the theorem prover KeYmaera X [23]. We encode the example shown in

Figure 7 (ϕ_{robust} in particular) as a $\text{d}\mathcal{L}$ formula $\pi(\xi(\phi_{\text{robust}}))$, which KeYmaera X is able to prove fully automatically.

Note that π *sequentially* composes the left and right programs, i.e., $\pi((\alpha, \beta)) = \alpha; \beta$. Though the encoding based on π is sound and complete, it may be challenging to directly verify certain $\text{d}\mathcal{L}_{\text{REL}}$ formulas with this encoding. We discuss next how to improve π by integrating techniques for both general relational reasoning and domain-specific relational reasoning.

5.2 Integrating Techniques in the Encoding

A REL instantiation like $\text{d}\mathcal{L}_{\text{REL}}$ naturally blends general relation constructs and domain-specific constructs, e.g., $\llbracket (x' = \theta \& \phi, y' = \delta \& \psi) \gg ([x]_l \leq [y]_r) \rrbracket$ expresses a *forall, exists* property on two differential equations. The integration opens great opportunities to utilize advances in both general relational reasoning and domain-specific relational reasoning. Using $\text{d}\mathcal{L}_{\text{REL}}$ as an example, we show how to improve the verification power by retrofitting the encoding π with various techniques.

In particular, we develop four variants of π , each with increasing verification power, achieved by incorporating additional *heuristics* in π . We either develop these heuristics or adapt them from prior work. We test these variants on a set of benchmarks, by (1) applying these variants on the benchmarks to generate $\text{d}\mathcal{L}$ formulas, and then (2) use the theorem prover KeYmaera X to verify the generated formulas. Figure 8 shows the benchmarks and the results.

Benchmarks. The benchmarks include a set of example properties from this paper and some existing work. In particular, we collect many motivating examples from [3], as it focuses on Kleene Algebra with Tests, which is directly relevant to dynamics logics. We are able to verify all examples from [3], except for two examples that require native support of integers (Ex.2.2) and arrays (Ex.2.3), which are not natively supported by KeYmaera X. In addition, we verify a few other representative and non-trivial case studies from existing work.

We elaborate on the four variants of π and their corresponding results. Note that, as shown in Figure 8, the default π can automatically verify three of these benchmarks.

Benchmarks \ Variants	Variants				
	π	π_{align}	$\pi_{\forall\exists}$	π_{ode}	π_{syn}
Fig.7 (Fig.8 in [48])	A	A	A	A	A
Ex.2.1 in [3]		S	S	S	S
Ex.2.4 in [3]	A	A	A	A	A
Ex.7.4 in [3]	A	A	A	A	A
Ex.7.5 in [3]			A	A	A
Ex.7.6 in [3]			A	A	A
Ex.1				A	A
Fig.8 in [49]				S	S
Cs.1 in [29]					UA
Cs.2 in [29]					UA

Fig. 8: Using variants of π to verify non-trivial relational properties. For notation, *Ex.* and *Cs.* respectively denote the *example properties* and *case studies* in the cited work. Table cells indicate how much automation KeYmaera X can verify the generated $\text{d}\mathcal{L}$ formulas. Notations A, S, and UA mean, respectively, fully automated, automated with less than 3 manual proof steps, and automated after user inputs.

π_{align} : π + **heuristics on alignment**. As shown by existing work [3, 41], verifying relational properties often requires effective *alignment* of programs and computations so that relational invariants can be easier specified. Consider the following property from existing work [3] written in $d\mathcal{L}_{REL}$:

$$[n]_i = [n]_r \rightarrow [i := n; r := 1] \llbracket (body^*, body^*); (?i = 0, ?i = 0) \rrbracket [r]_i = [r]_r$$

Where $body \equiv ?i \neq 0; r := r \times i; i := i - 1$. An effective alignment for verifying this property is to rewrite $(body^*, body^*)$ into $(body, body)^*$ which run the left and right executions in lockstep. We develop π_{align} that adds into π a heuristic that encodes a proof rule from existing work (Rule (3) from [3]) that soundly performs such a rewrite. Apply π_{align} to the property above produces a $d\mathcal{L}$ formula that can be verified with KeYmaera X with only one manual step.

$\pi_{\forall\exists}$: π_{align} + $\forall\exists$ **heuristics**. Certain $\forall\exists$ properties on loops can be reduced to $\forall\forall$ properties for easier verification. In particular, the following reasoning is sound, as the lockstep of $(\alpha, \beta)^*$ presents a witness for the $\forall\exists$ property:

$$\text{if } \phi_{\tau} \rightarrow \llbracket (\alpha, \beta)^* \rrbracket \psi_{\tau} \text{ and } \phi_{\tau} \rightarrow \langle\langle ?\top, \beta \rangle^* \rangle \top \text{ then } \phi_{\tau} \rightarrow \llbracket (\alpha^*, \beta^*) \rrbracket \psi_{\tau}$$

Here, the condition $\phi_{\tau} \rightarrow \langle\langle ?\top, \beta \rangle^* \rangle \top$ ensures β has valid executions from ϕ_{τ} .

This rule can be generalized to support reasoning with more program constructs, such as $\llbracket (\alpha_1; \alpha^*; \alpha_2, \beta_1; \beta^*; \beta_2) \rrbracket \psi_{\tau}$.

We develop $\pi_{\forall\exists}$ by adding these $\forall\exists$ rules above into π_{align} , and test it on benchmarks from existing work (Ex.7.5 and Ex.7.6 from [3]). KeYmaera X is able to verify the generated $d\mathcal{L}$ formulas automatically.

π_{ode} : $\pi_{\forall\exists}$ + **heuristics on dynamics**. In addition to integrating general relational reasoning (showcased in variants π_{align} and $\pi_{\forall\exists}$), we can integrate domain-specific relational reasoning. For example, consider the following example on a relationship between two dynamics of exponential decay:

$$[x]_i \geq [x]_r \wedge [x]_i \geq 0 \wedge A > 0 \wedge A \geq B \rightarrow \llbracket (x' = A \times x, x' = B \times x) \rrbracket [x]_i \geq [x]_r \quad (1)$$

This property cannot be verified with KeYmaera X using the variants above. However, this $\forall\exists$ property should hold since the x value of the right dynamic would always be lower if *the two dynamics evolve for the same duration of time*. We can derive heuristics based on this observation. Consider a formula of the form $\phi_{\tau} \rightarrow \llbracket (x' = \theta \& \phi_1, y' = \delta \& \phi_2) \rrbracket \psi_{\tau}$ (this formula has been properly renamed, so the left and right programs refer to disjoint sets of variables). After applying the π function, we get a $d\mathcal{L}$ formula of the form: $\phi \rightarrow [x' = \theta \& \phi_1][y' = \delta \& \phi_2] \psi$, which can be reasoned with the following heuristic:

$$\text{if } \phi \rightarrow [x' = \theta, y' = \delta \& (\phi_1 \wedge \phi_2)] \psi \text{ and } \phi \rightarrow [x' = \theta, y' = \delta](\phi_1 \rightarrow \phi_2)$$

$$\text{then } \phi \rightarrow [x' = \theta \& \phi_1][y' = \delta \& \phi_2] \psi$$

The first condition of this heuristic merges two dynamics into one and forces the two dynamics to evolve for the same duration of time. A formula of the form $\phi \rightarrow [x' = \theta, y' = \delta \& (\phi_1 \wedge \phi_2)] \psi$ is often easier to prove. The second condition on evolution constraints indicates there exists executions of the right dynamics if the left has valid executions. The second condition is needed to ensure *there exists an execution of β* for a $\llbracket (\alpha, \beta) \rrbracket \phi_{\tau}$ modality.

We develop π_{ode} by adding this heuristic into $\pi_{\forall\exists}$, and we manage to automatically verify EQ.1 in KeYmaera X. In addition to Eq.1, we apply π_{ode} to a non-trivial case study on a water tank from existing work [49]. Its main property

is a $\forall\exists$ one on a system-level model of the water tank. The function π_{ode} first applies the heuristic added in $\pi_{\forall\exists}$, and then applies the heuristic added in π_{ode} . We manage to verify the $d\mathcal{L}$ formula generated by π_{ode} on this property with only one manual proof step.

$\pi_{syn} \cdot \pi_{ode} + \text{User Inputs.}$ We further investigate the utility of the encoding by adding a heuristic derived from recent work on relational verification of $\forall\forall$ properties on dynamics (Thm.24[29]). The heuristic aims to reason with $d\mathcal{L}_{REL}$ formulas of the form: $\phi_{\tau} \rightarrow \llbracket (x' = \theta \& \phi_1, y' = \delta \& \phi_2) \rrbracket \psi_{\tau}$. However, instead of fixing the duration passed by the left and right dynamics as in π_{ode} , this heuristic asks for user inputs that would allow the two dynamics to *synchronize* at different time points, e.g., compare the values of x and y after the left and right dynamics evolve for different durations.

We develop a variant π_{syn} that asks users to provide inputs when it sees appropriate $d\mathcal{L}_{REL}$ formulas, and then produces corresponding $d\mathcal{L}$ formulas. We test π_{syn} with two case studies from [29]: after obtaining the correct inputs, π_{syn} produces $d\mathcal{L}$ formulas that can be automatically verified by KeYmaera X.

Discussion. Our experiments demonstrate the feasibility of integrating techniques into the encoding, thereby promoting automated verification of relational properties. Note that we manually develop all these variants of π . How to systematically and automatically integrate different techniques is an interesting and useful future work. The main difficulty is, given an input REL formula, to identify the heuristics to use in its transformation. We plan to develop algorithms to automate the selection of heuristics.

6 Related Work

BiKAT. The closest related work is the BiKAT [3], an extension of KAT to help reason with the alignments of relational properties through equational reasoning. A main focus of BiKAT is to ensure *adequacy*: the aligned program should capture all computations of the original program. Our biprogram construct is analogous to the design in BiKAT. However, a major difference between our work and BiKAT is applicability: the main focus of BiKAT is alignments of two programs, while we also concern domain-specific reasoning. The reasoning rules for BiKAT can be encoded in the translations (as shown in Section 5.2), but BiKAT cannot be used to reason with many $d\mathcal{L}_{REL}$ formulas. In addition, the design of modalities promotes expressing complex relational properties concisely, which is not a motivation behind BiKAT.

Relational Reasoning for Dynamic Logic. Beckert et al. have done a series of work on extending dynamic logic with trace modalities [10, 9, 26]. They further apply trace modalities to check secure information flow in the setting of concurrent programs [26]. Their work focus on first-order dynamic logic with deterministic programs. The limitation to deterministic programs is significant because nondeterminism is essential to the utility of dynamic logics, especially in AI-enabled systems [24]. Gutsfeld et al. introduce an expressive extension

to propositional dynamic logic to check hyperproperties by introducing path quantifiers [27]. Algorithms for model checking these hyperproperties are introduced. Tool support is left for future work. In comparison, the REL extension goes beyond propositional dynamic logic, and it can reuse existing tools to verify relational properties.

Relational Reasoning for Differential Dynamic Logic.. Various approaches have been proposed to analyze specific relational properties in $\text{d}\mathcal{L}$. A primitive is introduced to express a refinement relation between two hybrid programs [31]. An expressive modal logic based on $\text{d}\mathcal{L}$ has been introduced to reason with nondeducibility [13]. Neither work provides tool support. Kolčák et al. introduce a relational extension of $\text{d}\mathcal{L}$ that focuses on reasoning about $\forall\forall$ on two dynamics [29]. Our translation for $\text{d}\mathcal{L}_{\text{REL}}$ is able to integrate this work. Xiang et al. introduce a formal framework in the setting of $\text{d}\mathcal{L}$ for modeling and analyzing the robustness of cyber-physical systems under sensor attacks [48]. In our case study for $\text{d}\mathcal{L}$, we express and verify this robustness property with $\text{d}\mathcal{L}_{\text{REL}}$, which is more succinct than the original example.

Biprogram.. Our biprogram is an adaptation of the work of Pottier and Simonet on information flow analysis for ML [38]. The work introduces an extension of ML that encodes a pair of ML terms as a single bracket construct. The bracket constructs cannot be nested. This is analogous to our design in that the biprogram contains two programs of the host logics. This bracket-based approach has been often used in proving noninterference, e.g., [51, 4]. The name “biprogram” is also used in other recent work on relational reasoning [6, 3].

Relational Program Logic. Another approach for relational verification is to use relational program logic [11, 50, 32]. Benton introduced Relational Hoare Logic for verifying program transformations [11], which provides a general framework for relational correctness proofs. Recent work proposes a notion to evaluate the design of Relational Hoare Logic [32]. Sousa and Dillig introduce a program logic, named Cartesian Hoare Logic (CHL), for verifying k-safety properties [42]. The work has been further extended for proving the correctness of 3-way merge [43]. They introduce a generalized form of Hoare triples to express relations between different program executions, and use SMT solvers to determine their satisfiability. D’Osualdo et al. [21] describe a logic for hyper-triple composition (LHC) based on weakest preconditions that can decompose a hypersafety proof along the boundary of hyper tuples, offering ways of combining multiple k-safety proofs. These relational program logics often are based on weakest pre-condition and only support $\forall\forall$ pre/post k-safety properties. In contrast, the REL extensions naturally support properties with mixed modalities. Moreover, the translations for REL extensions open opportunities in combining general relational reasoning and domain-specific reasoning. We believe these relational program logics can be also integrated as a part of the translations, which we leave to future work.

Self-Composition and Product Program.. The encodings introduced in this work are inspired by self-composition, which is a common approach for relational reasoning [8, 19, 7, 16]. Approaches based on self-composition are often syntax-

directed, that is, they compose two programs that have similar, if not identical, syntactic constructs. Alternatively, property directed self-composition [41] tackles programs that cannot be easily aligned in a composition. It composes programs (or copies of the same program) by finding good alignment between the copies in order to have expressible assertions. The encodings in this work can be viewed as special kinds of self-compositions that shall be adapted for different host logics. The encoding can benefit from advances in self-composition, as we have shown in Section 5.2.

Other Related Work. Noninterference is a well-known relational property for secure information flow and has been widely studied, e.g., [38, 44, 47, 5, 15]. REL extensions bring in several major benefits for properties like noninterference. First, noninterference reasons about program executions, and modalities in REL capture executions intuitively. Different variants of noninterference can be concisely expressed, as shown in the examples of Section 3. Second, the notion of noninterference can be easily extended to different domains (host logics), as the REL extension connects the common program constructs with domain-specific constructs. And third, verifying noninterference can leverage more existing tools, as the translations can integrate different techniques.

Relational verification has been studied in the setting of temporal-style logics, such as HyperLTL and HyperCTL, have been introduced to model relational properties [17, 12, 46, 27]. Model checking is often used to verify these properties. In contrast, REL extensions are more expressive on program behaviors than logics like HyperLTL and HyperCTL. Translations for the extensions may leverage the model checkers for HyperLTL or HyperCTL if the correct heuristics are developed, which we leave to future work.

7 Conclusion

We introduce a general and lightweight relational extension for dynamic logics. We expect that the extension can be instantiated for almost any dynamic logic, in the same way that one would expect that almost any dynamic logic could be extended (syntactically and semantically) with first-order operators. Verifying REL formulas can leverage existing tools developed for the host logics. Moreover, the instantiations show great potential in promoting verification by integrating general relational reasoning and domain-specific relational reasoning.

Future work. One immediate next step is to *systematically* integrate techniques for general relational reasoning and domain-specific relational reasoning into the encoding. Proof automation with the translations is an important and interesting future work. Additionally, we believe the natural expression of relational properties by the REL extension opens new opportunities to explore more powerful verification techniques. We plan to develop a novel proof system for the REL extension that allows us to prove relational properties that cannot be proven by existing tools and techniques.

References

1. Ahrendt, W., Beckert, B., Bubel, R., Hähnle, R., Schmitt, P.H., Ulbrich, M.: *Deductive Software Verification—The KeY Book: From Theory to Practice*, vol. 10001. Springer (2016)
2. Allen, P.: A comparison of non-interference and non-deducibility using CSP. In: CSF. pp. 43–44 (1991)
3. Antonopoulos, T., Koskinen, E., Le, T.C., Nagasamudram, R., Naumann, D.A., Ngo, M.: An algebra of alignment for relational verification. *PLDI* **7**, 573–603 (2023)
4. Arden, O., Myers, A.C.: A calculus for flow-limited authorization. In: CSF. pp. 135–149 (2016)
5. Balliu, M., Dam, M., Le Guernic, G.: Epistemic temporal logic for information flow security. In: PLAS. pp. 1–12 (2011)
6. Banerjee, A., Nagasamudram, R., Naumann, D.A., Nikouei, M.: A relational program logic with data abstraction and dynamic framing. *arXiv preprint arXiv:1910.14560* (2019)
7. Barthe, G., Crespo, J.M., Kunz, C.: Relational verification using product programs. In: FM. pp. 200–214 (2011)
8. Barthe, G., D’Argenio, P.R., Rezk, T.: Secure information flow by self-composition. In: CSF. pp. 100–114 (2004)
9. Beckert, B., Bruns, D.: Dynamic logic with trace semantics. In: CADE. pp. 315–329 (2013)
10. Beckert, B., Schlager, S.: A sequent calculus for first-order dynamic logic with trace modalities. In: IJCAR. pp. 626–641 (2001)
11. Benton, N.: Simple relational correctness proofs for static analyses and program transformations. In: POPL. pp. 14–25 (2004)
12. Beutner, R., Finkbeiner, B.: Autohyper: Explicit-state model checking for Hyper-LTL. In: TACAS. pp. 145–163 (2023)
13. Bohrer, B., Platzer, A.: A hybrid, dynamic logic for hybrid-dynamic information flow. In: LICS. pp. 115–124 (2018)
14. Chong, S., Lanotte, R., Merro, M., Tini, S., Xiang, J.: Quantitative robustness analysis of sensor attacks on cyber-physical systems. In: HSCC. pp. 1–12 (2023)
15. Chudnov, A., Naumann, D.A.: Assuming you know: Epistemic semantics of relational annotations for expressive flow policies. In: CSF. pp. 189–203 (2018)
16. Churchill, B., Padon, O., Sharma, R., Aiken, A.: Semantic program alignment for equivalence checking. In: PLDI. pp. 1027–1040 (2019)
17. Clarkson, M.R., Finkbeiner, B., Koleini, M., Micinski, K.K., Rabe, M.N., Sánchez, C.: Temporal logics for hyperproperties. In: POST. pp. 265–284 (2014)
18. Clarkson, M.R., Schneider, F.B.: Hyperproperties. *Journal of Computer Security* **18**(6), 1157–1210 (2010)
19. Darvas, Á., Hähnle, R., Sands, D.: A theorem proving approach to analysis of secure information flow. In: PerCom. pp. 193–209 (2005)
20. De Giacomo, G., Vardi, M.Y.: Linear temporal logic and linear dynamic logic on finite traces. In: IJCAI (2013)
21. D’Osualdo, E., Farzan, A., Dreyer, D.: Proving hypersafety compositionally. In: OOPSLA2. pp. 289–314 (2022)
22. Francez, N.: Product properties and their direct verification. *Acta informatica* **20**(4), 329–344 (1983)

23. Fulton, N., Mitsch, S., Quesel, J.D., Völz, M., Platzer, A.: KeYmaera X: An axiomatic tactical theorem prover for hybrid systems. In: CADE. pp. 527–538 (2015)
24. Fulton, N., Platzer, A.: Verifiably safe off-model reinforcement learning. In: TACAS. pp. 413–430 (2019)
25. Goguen, J.A., Meseguer, J.: Security policies and security models. In: S&P. pp. 11–20 (1982)
26. Grahl, D.: Deductive Verification of Concurrent Programs and its Application to Secure Information Flow for Java. Ph.D. thesis, Karlsruhe Institute of Technology (29 Oct 2015)
27. Gutsfeld, J.O., Müller-Olm, M., Ohrem, C.: Propositional dynamic logic for hyperproperties. In: CONCUR (2020)
28. Harel, D., Kozen, D., Tiuryn, J.: Dynamic Logic. MIT Press (2000)
29. Kolčák, J., Dubut, J., Hasuo, I., Katsumata, S.y., Sprunger, D., Yamada, A.: Relational differential dynamic logic. In: TACAS. pp. 191–208 (2020)
30. Kozen, D.: Kleene algebra with tests. TOPLAS **19**(3), 427–443 (1997)
31. Loos, S.M., Platzer, A.: Differential refinement logic. In: LICS. pp. 505–514 (2016)
32. Nagasamudram, R., Naumann, D.A.: Alignment completeness for relational hoare logics. In: LICS. pp. 1–13 (2021)
33. Naumann, D.A.: Thirty-seven years of relational hoare logic: remarks on its principles and history. In: ISoLA. pp. 93–116 (2020)
34. Platzer, A.: Differential dynamic logic for hybrid systems. Journal of Automated Reasoning **41**(2), 143–189 (2008)
35. Platzer, A.: The complete proof theory of hybrid systems. In: LICS. pp. 541–550 (2012)
36. Platzer, A.: A complete uniform substitution calculus for differential dynamic logic. Journal of Automated Reasoning **59**(2), 219–265 (2017)
37. Platzer, A.: Logical foundations of cyber-physical systems, vol. 662. Springer (2018)
38. Pottier, F., Simonet, V.: Information flow inference for ML. In: POPL. pp. 319–330 (2002)
39. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. Journal on Selected Areas in Communications **21**(1), 5–19 (2003)
40. Satake, Y., Unno, H.: Propositional dynamic logic for higher-order functional programs. In: CAV. pp. 105–123 (2018)
41. Shemer, R., Gurfinkel, A., Shoham, S., Vizel, Y.: Property directed self composition. In: CAV. pp. 161–179 (2019)
42. Sousa, M., Dillig, I.: Cartesian hoare logic for verifying k-safety properties. In: PLDI. pp. 57–69 (2016)
43. Sousa, M., Dillig, I., Lahiri, S.K.: Verified three-way program merge. In: OOPSLA. pp. 145–164 (2016)
44. Stefan, D., Russo, A., Mitchell, J.C., Mazières, D.: Flexible dynamic information flow control in Haskell. In: Haskell. pp. 95–106 (2011)
45. Terauchi, T., Aiken, A.: Secure information flow as a safety problem. In: SAS. pp. 352–367 (2005)
46. Wang, Y., Nalluri, S., Pajic, M.: Hyperproperties for robotics: Planning via HyperLTL. In: ICRA. pp. 8462–8468 (2020)
47. Xiang, J., Chong, S.: Co-Inflow: coarse-grained information flow control for Java-like languages. In: S&P (2021)

- 48. Xiang, J., Fulton, N., Chong, S.: Relational analysis of sensor attacks on Cyber-Physical Systems. In: CSF (2021)
- 49. Xiang, J., Lanotte, R., Tini, S., Chong, S., Merro, M.: Measuring robustness in cyber-physical systems under sensor attacks. *Nonlinear Analysis: Hybrid Systems* **56** (2025)
- 50. Yang, H.: Relational separation logic. *Theoretical Computer Science* **375**(1-3), 308–334 (2007)
- 51. Zheng, L., Myers, A.C.: Dynamic security labels and static information flow control. *International Journal of Information Security* **6**(2-3), 67–84 (2007)