

Verification of Data-Aware Processes

Verification Logics

Diego Calvanese, Marco Montali

Research Centre for Knowledge and Data (KRDB)
Free University of Bozen-Bolzano, Italy



29th European Summer School in Logic, Language, and Information
(ESLLI 2017)

Toulouse, France – 17–28 July 2017

Outline

- 1 Verification Logics
- 2 First-order μ -Calculus
- 3 History and Persistence-Preserving μ -Calculus
- 4 First-order Linear Temporal Logics

Verification logics

Temporal/dynamic logics expressing (un)desired properties about the dynamic system of interest.

- **Branching-time logics** are directly interpreted over the transition system.
Notable examples: CTL, μ -calculus
- **Linear-time logics** are interpreted over the runs represented by the transition system. A property holds over the transition system if it is true in every run.
Notable example: LTL

We consider in particular LTL and μ -calculus ($\mu\mathcal{L}$) as two representative verification logics.

The verification logic $\mu\mathcal{L}$

Remember that $\mu\mathcal{L}$:

- Is equipped with two local temporal operators to move over one ($\langle\langle-\rangle\rangle$) or all ($[[-]]$) successors of the current state.
- Employs fixpoint constructs to express sophisticated properties defined via induction or co-induction.
- Subsumes virtually all propositional temporal/dynamic logics, such as PDL, LTL, CTL, CTL*.

Note: The encoding into $\mu\mathcal{L}$ of linear-time properties incurs in an exponential blow-up.

Verification logics for DCDSs

Propositional temporal logics do not suffice!

We need **first-order temporal logics**:

- To inspect **data**: **FO queries**
- To capture system **dynamics**: **temporal modalities**
- To track the **evolution of objects**: FO **quantification across** states

Example:

It is **always** the case that **every order** is **eventually** either **cancelled** or **paid**.

$$\mathbf{G}(\forall x. \mathbf{Order}(x) \rightarrow \mathbf{F}(\mathbf{State}(x, \mathbf{cancelled}) \vee \mathbf{State}(x, \mathbf{paid})))$$

Outline

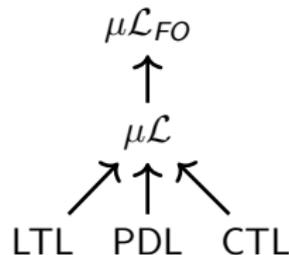
- 1 Verification Logics
- 2 First-order μ -Calculus**
- 3 History and Persistence-Preserving μ -Calculus
- 4 First-order Linear Temporal Logics

First-order μ -calculi for DCDSs

We employ variants of **first-order μ -calculus** ($\mu\mathcal{L}_{FO}$):

$$\Phi ::= Q \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x.\Phi \mid \langle - \rangle\Phi \mid Z \mid \mu Z.\Phi$$

- Extends the propositional μ -calculus $\mu\mathcal{L}$ with first-order quantification.
- The first-order quantifiers range over all objects in the transition system (and not only over those in the current state or in the current run).



We also adopt the standard abbreviations, including:

- $[-]\Phi$ for $\neg\langle - \rangle\neg\Phi$
- $\nu Z.\Phi$ for $\neg\mu Z.\Phi_{[Z/\neg Z]}$

Example

$$\forall x.\text{Student}(x) \rightarrow \mu Z.((\exists y.\text{Graduate}(x, y)) \vee \langle - \rangle Z)$$

For each student x (in the current state), there exists an evolution that eventually leads to the graduation of x (with some final mark y).

Semantics of $\mu\mathcal{L}_{FO}$

$\mu\mathcal{L}_{FO}$ formulae: interpreted over RTS $\Upsilon = \langle \Delta, \mathcal{R}, S, s_0, db, \Rightarrow \rangle$, using valuations:

- v : **individual variable valuation** mapping each individual variable x to a value in Δ .
- V : **predicate variable valuation**, parameterized by v , and mapping each predicate variable Z to a subset $V(v, Z)$ of S .

Evaluation of a $\mu\mathcal{L}_{FO}$ formula Φ over Υ : given by an **extension function** $(\Phi)_{(v,V)}^{\Upsilon}$ mapping Φ to the set of states in S where Φ holds.

$$(\varphi)_{(v,V)}^{\Upsilon} = \{s \mid s \in S \text{ and } \langle db(s), v \rangle \models \varphi\}$$

$$(\neg\Phi)_{(v,V)}^{\Upsilon} = S \setminus (\Phi)_{(v,V)}^{\Upsilon}$$

$$(\Phi_1 \wedge \Phi_2)_{(v,V)}^{\Upsilon} = (\Phi_1)_{(v,V)}^{\Upsilon} \cap (\Phi_2)_{(v,V)}^{\Upsilon}$$

$$(\exists x.\Phi)_{(v,V)}^{\Upsilon} = \{s \in S \mid \exists d \in \Delta. s \in (\Phi)_{(v,V)[x/d]}^{\Upsilon}\}$$

stands for (v', V) ,
where v' is as v
except that $v'(x) = d$

$$(\langle \rightarrow \rangle \Phi)_{(v,V)}^{\Upsilon} = \{s \in S \mid \exists s' \in S. s \Rightarrow s' \text{ and } s' \in (\Phi)_{(v,V)}^{\Upsilon}\}$$

$$(Z)_{(v,V)}^{\Upsilon} = V(v, Z)$$

$$(\mu Z.\Phi)_{(v,V)}^{\Upsilon} = \bigcap \{\mathcal{E} \subseteq S \mid (\Phi)_{(v,V)[Z/\mathcal{E}]}^{\Upsilon} \subseteq \mathcal{E}\}$$

stands for (v, V') ,
where V' is as V
except that $V'(v, Z) = \mathcal{E}$

Model checking $\mu\mathcal{L}_{FO}$

When the $\mu\mathcal{L}_{FO}$ formula Φ is **closed**, $(\Phi)_{(v,V)}^{\Upsilon}$ does not depend on the valuations v and V , and we simply denote it as $(\Phi)^{\Upsilon}$.

Model checking an RTS

Input:

- an RTS $\Upsilon = \langle \Delta, \mathcal{R}, S, s_0, db, \Rightarrow \rangle$
- a closed $\mu\mathcal{L}_{FO}$ formula Φ

Output: yes, iff $s_0 \in (\Phi)^{\Upsilon}$

In this case, we write $\Upsilon \models \Phi$.

Model checking a DCDS

Input:

- a DCDS \mathcal{X}
- a closed $\mu\mathcal{L}_{FO}$ formula Φ

Output: yes, iff $\Upsilon_{\mathcal{X}} \models \Phi$

In this case, we write $\mathcal{X} \models \Phi$.

Outline

- 1 Verification Logics
- 2 First-order μ -Calculus
- 3 History and Persistence-Preserving μ -Calculus**
- 4 First-order Linear Temporal Logics

History-preserving μ -calculus ($\mu\mathcal{L}_A$)

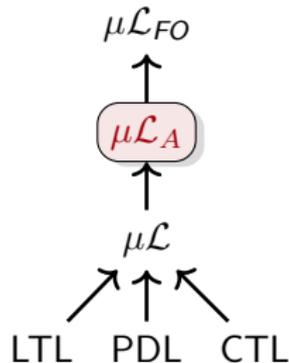
Active-domain quantification: restricted to those individuals *present in the current database*.

$$\exists x.\Phi \quad \rightsquigarrow \quad \exists x.\text{LIVE}(x) \wedge \Phi$$

$$\forall x.\Phi \quad \rightsquigarrow \quad \forall x.\text{LIVE}(x) \rightarrow \Phi$$

where $\text{LIVE}(x)$ states that x is present in the current active domain (easily expressible in FO).

Note: $\mu\mathcal{L}_A$ is a syntactic restriction of $\mu\mathcal{L}_{FO}$.



Example

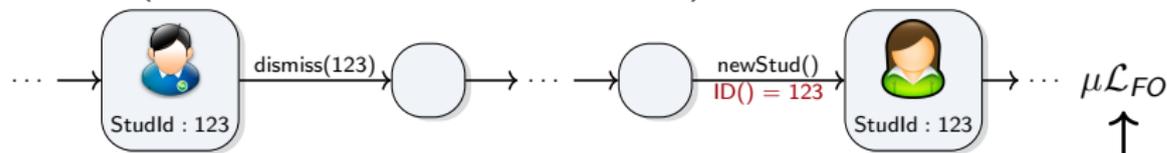
$$\nu W.(\forall x.\text{LIVE}(x) \wedge \text{Student}(x) \rightarrow \mu Z.(\exists y.\text{LIVE}(y) \wedge \text{Graduate}(x, y) \vee \langle - \rangle Z) \wedge [-]W)$$

Along every path, it is always true, for each student x , that there exists an evolution eventually leading to a graduation of the student (with some final mark y).

Note: No guarantee that all such students graduate within the same run.

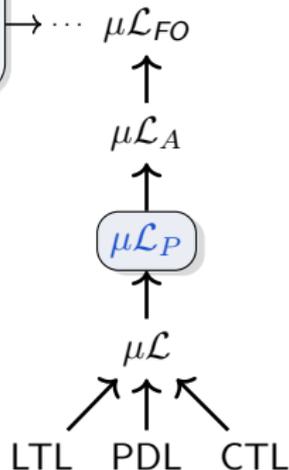
Persistence-preserving μ -calculus ($\mu\mathcal{L}_P$)

In some cases, objects maintain their identity only if they **persist** in the active domain (cf. business artifacts and their IDs).



$\mu\mathcal{L}_P$ restricts $\mu\mathcal{L}_A$ to **quantification over persisting objects only**, i.e., objects that *continue* to be LIVE.

$$\begin{aligned} \exists x.\Phi &\sim \exists x.\text{LIVE}(x) \wedge \Phi \\ \forall x.\Phi &\sim \forall x.\text{LIVE}(x) \rightarrow \Phi \\ \langle - \rangle \Phi(\vec{x}) &\sim \begin{cases} \text{LIVE}(\vec{x}) \wedge \langle - \rangle \Phi(\vec{x}) & \text{(strong persistence)} \\ \text{LIVE}(\vec{x}) \rightarrow \langle - \rangle \Phi(\vec{x}) & \text{(weak persistence)} \end{cases} \\ [-] \Phi(\vec{x}) &\sim \begin{cases} \text{LIVE}(\vec{x}) \wedge [-] \Phi(\vec{x}) & \text{(strong persistence)} \\ \text{LIVE}(\vec{x}) \rightarrow [-] \Phi(\vec{x}) & \text{(weak persistence)} \end{cases} \end{aligned}$$



Note: $\mu\mathcal{L}_P$ is a syntactic restriction of $\mu\mathcal{L}_A$.

Strong vs. weak persistence

Strong persistence: property falsified by an object that disappears

$$\nu W.(\forall x.\text{LIVE}(x) \wedge \text{Student}(x) \rightarrow \\ \mu Z.(\exists y.\text{LIVE}(y) \wedge \text{Graduate}(x, y) \vee (\text{LIVE}(x) \wedge \langle - \rangle Z)) \wedge [-]W)$$

Along every path, it is always true, for each student x , that there exists an evolution in which x **persists in the database until** she eventually graduates.

Weak persistence: property verified by an object that disappears

$$\nu W.(\forall x.\text{LIVE}(x) \wedge \text{Student}(x) \rightarrow \\ \mu Z.(\exists y.\text{LIVE}(y) \wedge \text{Graduate}(x, y) \vee (\text{LIVE}(x) \rightarrow \langle - \rangle Z)) \wedge [-]W)$$

Along every path, it is always true, for each student x , that there exists an evolution in which **either x does not persist, or** she eventually graduates.

Outline

- 1 Verification Logics
- 2 First-order μ -Calculus
- 3 History and Persistence-Preserving μ -Calculus
- 4 First-order Linear Temporal Logics**

First-order linear temporal logics for DCDSs

LTL-FO extends propositional LTL with the possibility of querying the system states using first-order formulas with quantification across:

$$\Phi ::= \varphi \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x.\Phi \mid \mathbf{X}\Phi \mid \Phi_1 \mathbf{U}\Phi_2$$

We also adopt the standard abbreviations, including:

- $\mathbf{F}\Phi$ for $\text{true} \mathbf{U}\Phi$ (Φ holds in the future)
- $\mathbf{G}\Phi$ for $\neg\mathbf{F}\neg\Phi$ (Φ holds globally)

Example

$$\forall x.\text{Student}(x) \rightarrow \mathbf{F}\exists y.\text{Graduate}(x, y)$$

For each student x (in the current state), x will graduate sometimes in the future (with some final mark y).

Note: all encountered students graduate within the same run.

Semantics of LTL-FO

Formulae of LTL-FO_A are interpreted over the infinite runs of a given **serial** RTS $\Upsilon = \langle \mathcal{R}, S, s_0, db, \Rightarrow \rangle$.

- **Serial RTS**: every state has at least one successor state.

An **(infinite) run** τ over Υ

is an infinite sequence $s_0 s_1 \dots$ of states in S , where the first state of the sequence corresponds to the initial state of Υ , and for every $i \in \mathbb{N}$, it is true that $s_i \Rightarrow s_{i+1}$. Given $j \in \mathbb{N}$, by $\tau(j)$ we denote the j -th state s_j of τ .

We inductively define when τ satisfies an LTL-FO_A formula Φ at position i under v , written $\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi$

$\langle \tau, i, v \rangle \models_{\text{LTL}} \varphi$	if $\langle db(\tau(i)), v \rangle \models \varphi$
$\langle \tau, i, v \rangle \models_{\text{LTL}} \neg \Phi$	if it is not the case that $\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi$
$\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi_1 \wedge \Phi_2$	if $\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi_1$ and $\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi_2$
$\langle \tau, i, v \rangle \models_{\text{LTL}} \exists x. \Phi$	if there exists $d \in \Delta$ such that $\langle \tau, i, v[x/d] \rangle \models_{\text{LTL}} \Phi$
$\langle \tau, i, v \rangle \models_{\text{LTL}} \mathbf{X} \Phi$	if $\langle \tau, i + 1, v \rangle \models_{\text{LTL}} \Phi$
$\langle \tau, i, v \rangle \models_{\text{LTL}} \Phi_1 \mathbf{U} \Phi_2$	if there exists $k \geq i$ such that $\langle \tau, k, v \rangle \models_{\text{LTL}} \Phi_2$ and for every j , if $i \leq j < k$ then $\langle \tau, j, v \rangle \models_{\text{LTL}} \Phi_1$

Model checking LTL-FO

When the LTL-FO formula Φ is **closed**, the satisfaction relation does not depend on the valuation v , and we simply denote satisfaction as $\langle \tau, i \rangle \models_{\text{LTL}} \Phi$.

LTL model checking an RTS

Input:

- an RTS $\Upsilon = \langle \Delta, \mathcal{R}, S, s_0, db, \Rightarrow \rangle$
- a closed LTL-FO formula Φ

Output: yes, iff for every run τ over Υ , $\langle \tau, 0 \rangle \models_{\text{LTL}} \Phi$.

In this case, we write $\Upsilon \models_{\text{LTL}} \Phi$.

LTL Model checking a DCDS

Input:

- a DCDS \mathcal{X}
- a closed LTL-FO formula Φ

Output: yes, iff $\Upsilon_{\mathcal{X}} \models_{\text{LTL}} \Phi$.

In this case, we write $\mathcal{X} \models_{\text{LTL}} \Phi$.

First-order LTL with restricted quantification

History-preserving quantification: **LTL-FO_A**

FO quantification ranges over current active domain only:

$$\exists x.\Phi \rightsquigarrow \exists x.\text{LIVE}(x) \wedge \Phi$$

$$\forall x.\Phi \rightsquigarrow \forall x.\text{LIVE}(x) \rightarrow \Phi$$

Example: $\forall x.\text{LIVE}(x) \wedge \text{Customer}(x) \rightarrow \mathbf{F} \text{Gold}(x)$

Persistence-preserving quantification: **LTL-FO_P**

FO quantification ranges over persisting individuals only.

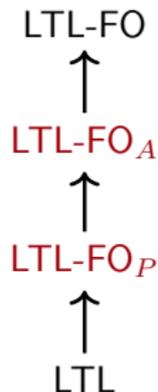
$$\exists x.\Phi \rightsquigarrow \exists x.\text{LIVE}(x) \wedge \Phi$$

$$\forall x.\Phi \rightsquigarrow \forall x.\text{LIVE}(x) \rightarrow \Phi$$

$$\mathbf{X} \Phi(\vec{x}) \rightsquigarrow \begin{cases} \text{LIVE}(\vec{x}) \wedge \mathbf{X} \Phi(\vec{x}) & \text{(strong persistence)} \\ \text{LIVE}(\vec{x}) \rightarrow \mathbf{X} \Phi(\vec{x}) & \text{(weak persistence)} \end{cases}$$

$$\Phi_1 \mathbf{U} \Phi_2(\vec{x}) \rightsquigarrow \begin{cases} (\text{LIVE}(\vec{x}) \wedge \Phi_1) \mathbf{U} \Phi_2(\vec{x}) & \text{(s.p.)} \\ (\text{LIVE}(\vec{x}) \wedge \Phi_1) \mathbf{U} (\text{LIVE}(\vec{x}) \rightarrow \Phi_2(\vec{x})) & \text{(w.p.)} \end{cases}$$

Example: $\forall x.(\text{LIVE}(x) \wedge \text{Gold}(x)) \rightarrow \neg(\text{LIVE}(x) \mathbf{U} \neg \text{Gold}(x))$



LTL with persistence-preserving quantification – Example

Consider:

$$\forall x. \text{Gold}(x) \rightarrow \mathbf{G} \text{Gold}(x)$$

$$\forall x. \text{Gold}(x) \rightarrow \neg \mathbf{F} \neg \text{Gold}(x)$$

$$\forall x. \text{Gold}(x) \rightarrow \neg(\text{true} \mathbf{U} \neg \text{Gold}(x))$$

With **strong persistence**:

$$\forall x. \text{LIVE}(x) \rightarrow (\text{Gold}(x) \rightarrow \neg(\text{LIVE}(x) \mathbf{U} \neg \text{Gold}(x)))$$

$$\forall x. (\text{LIVE}(x) \wedge \text{Gold}(x)) \rightarrow \neg(\text{LIVE}(x) \mathbf{U} \neg \text{Gold}(x))$$

With **weak persistence**:

$$\forall x. \text{LIVE}(x) \rightarrow (\text{Gold}(x) \rightarrow \neg(\text{LIVE}(x) \mathbf{U}(\text{LIVE}(x) \rightarrow \neg \text{Gold}(x))))$$

$$\forall x. (\text{LIVE}(x) \wedge \text{Gold}(x)) \rightarrow \neg(\text{LIVE}(x) \mathbf{U}(\text{LIVE}(x) \rightarrow \neg \text{Gold}(x)))$$