

Model Checking Reo Connectors with mCRL2

Natallia Kokash

Introduction

- COMPAS Project
 - Formalization of Business Process Models
 - Formalization of compliance requirements
- Reo
- mCRL2
- Model checking Reo process models with mCRL2
 - Translation from Reo to mCRL2
 - Unified Control and Data Flow Analysis
- Tool support and examples
- Conclusions and Future Work

COMPAS project

- **COMPAS**

- Compliance-driven Models, Languages, and Architectures for Services

- **Goal**

- Ensure dynamic and on-going **compliance** of software **services** to business regulations and user requirements

- **Methodology, terminology and research results**

- <http://www.compas-ict.eu/>
- *Compliance* is conformity in fulfilling compliance requirements
- *Compliance requirement* is a constraint or assertion that results from the interpretation of the compliance sources
- *Compliance source* is a document that is the origin of compliance requirements (e.g., SOX, HIPAA, license)

Formal BP analysis in COMPAS

- *Business process* is a composition of activities into a structured order that implements the procedure to be followed in order to achieve a business goal
- *Behavioral model* is a description of how an actor (e.g., stakeholder, service) acts or interacts with other actors
- *Compliance rule* is an operative definition of a compliance requirement
- **Goal: Check whether compliance rules hold for formal behavioral models of business processes**
- *Kokash, N., Arbab, F.: "Formal Behavioral Modeling and Compliance Analysis for Service-Oriented Systems", FMCO'08, vol. 5751 of LNCS, Springer, pp. 21-41*

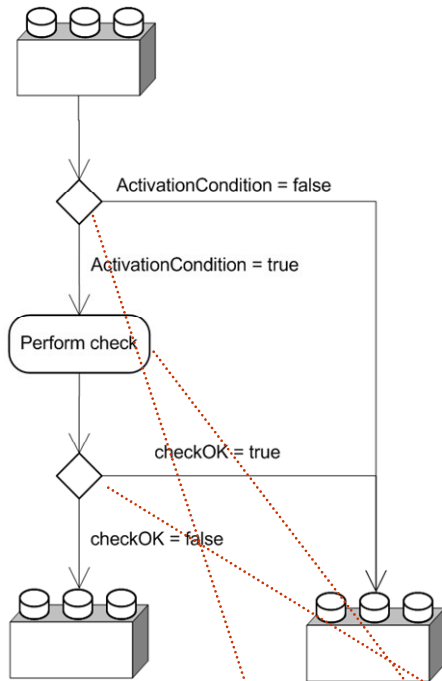
Reo coordination language

- Components (services) communicate through channels composed to complex connectors
- Semantics:
 - Constraint automata
 - Intentional automata (support for context-dependency)

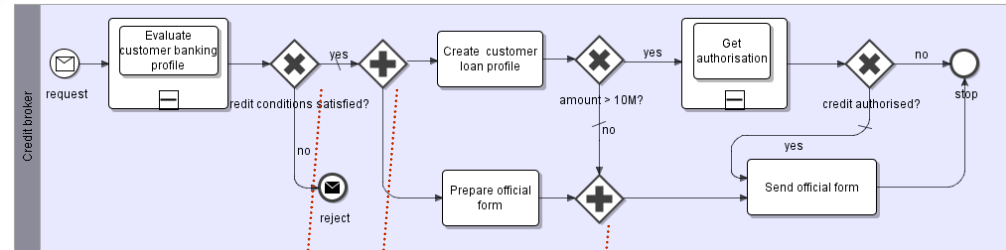
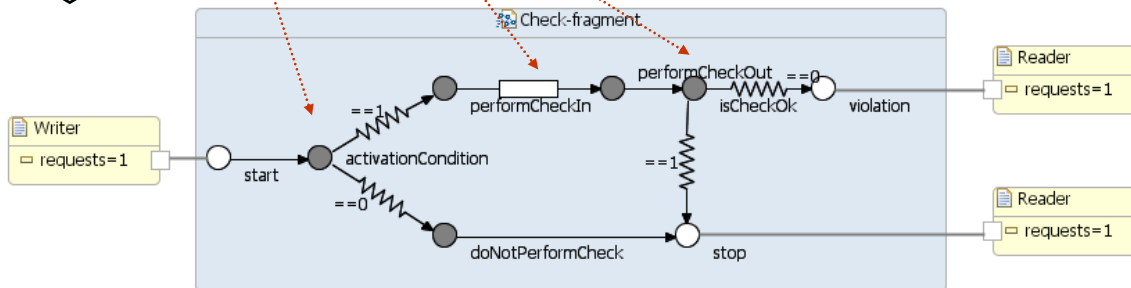
Sync, SyncDrain	LossySync	AsyncDrain	Fifo1

Replicator	Merger	Router

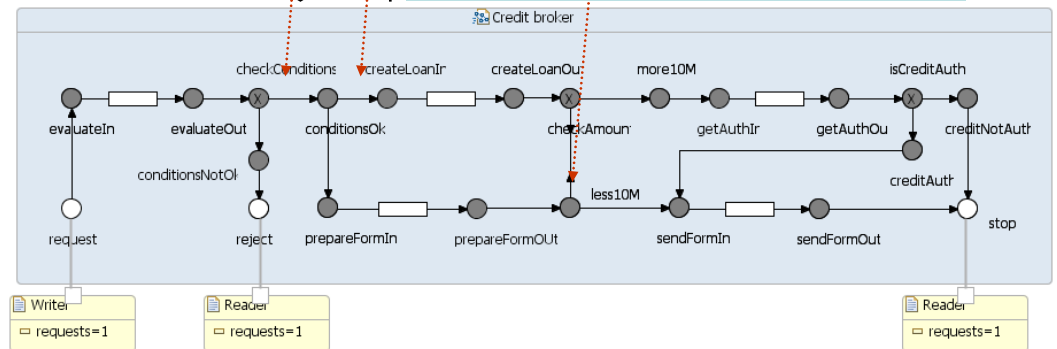
Formalization of business processes



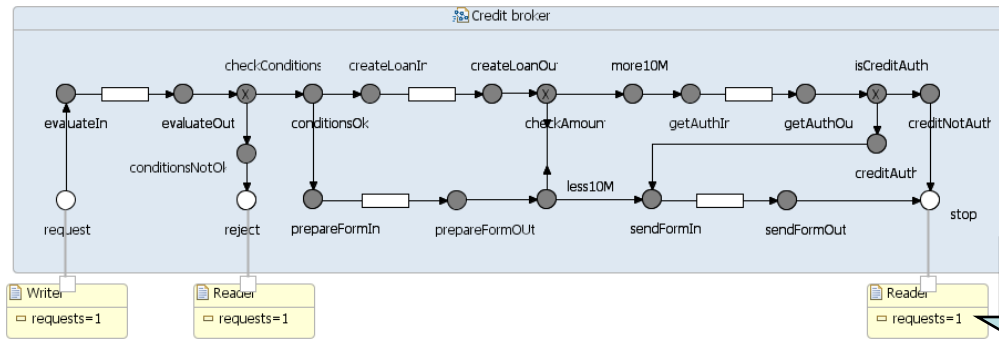
UMLADs2Reo converter




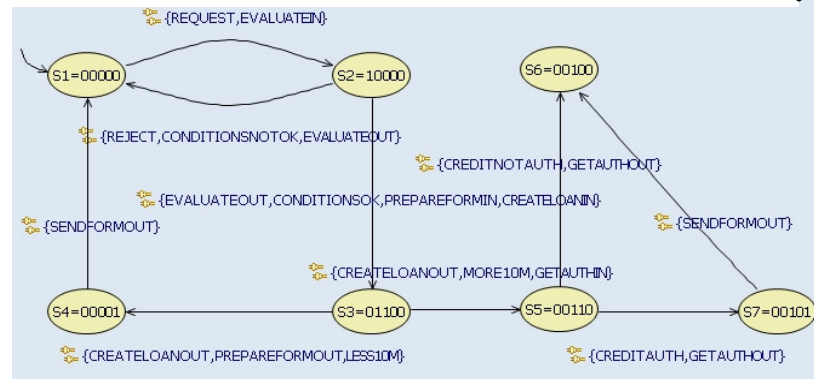
BPMN2Reo converter



Formalization of business processes



Reo2ExtendedAutomata
converters



- Check that admissible states *reject* or *sendFormOut* are eventually reached
 - $\mathbf{G}(\text{request} \rightarrow \mathbf{F}(\text{reject} \cup \text{sendFormOut}))$

Business process analysis

- Control flow analysis
 - *Deadlock* is a situation when a service or a process stays idle waiting for resources permanently blocked by another party or messages that will never arrive.
 - *Livelock (starvation of progress)* is a situation when a process or a service is not blocked but doesn't make any progress nevertheless.
 - *Problems with synchronization* conflicts can be introduced by joining fork concurrent paths with a merge structure which in some workflow specification languages (e.g., BPMN) results into unintentional multiple activation of states that follow the merge node.
- Data flow analysis
 - Deadlocks, livelocks and problems with synchronization may appear due to improper conditions on data (e.g., while data-based choice gateway is used in BPMN).
 - Compliance requirements are data dependent.

mCRL2

- Behavioral specification language
- Associated toolset
 - Developed at TU Eindhoven (+ LaQuSo, CWI and Twente University)
- Based on the algebra of communicating processes (ACP)
- Extended with data and time
 - Built-in data types: *Bool, Nat, Pos, Int, Real*
 - Algebraic data types
 - constructors, recognition and projection functions
 - Built-in support for lists, sets and bags
 - User-defined functions (λ calculus)
- Number of industrial case studies
- <http://www.mcrl2.org/>

mCRL2 specification language

- **Actions** are atomic events (e.g. a firing of a port or a request arrival in a Reo connector)
- **Processes** are the active entities defined as expressions over actions and other processes
 - **Multiaction**: a/b (synchronized actions)
 - **Alternative composition**: $a + b$ (nondeterministic choice)
 - **Sequence composition**: $a.b$ (b started after a)
 - **Conditional**: $exp \rightarrow a \diamond b$ (if-then-else)
 - **At operator**: $a^c t$ (action a happens at time t)
 - **Parallel composition**: $a//b$ (interleavings $a.b + b.a + a|b$)
- Actions and processes can be parametrized with **data**
 - **Summation**: $\sum_{d \in D} a(d)$ ($a(d_1) + a(d_2) + a(d_3) \dots$)

mCRL2 specification language

- **Renaming:** $\rho_R(a)$ where R is a set of renamings of the form $b \rightarrow c$, meaning that every occurrence of b in a is replaced by c
- **Hiding:** $\tau_H(a)$ renames all actions of H in a to τ
- **Restriction (allow):** $\nabla_R(a)$ where R specifies which actions are allowed to occur in a
- **Blocking:** $\partial_B(a)$ where B is a set of actions that is not allowed to occur in a
- **Communication:** $\Gamma_C(p)$, where C is a set of allowed communications of the form $a_0/\dots/a_n \rightarrow c$, $n \geq 1$ which means that every group of actions $a_0/\dots/a_n$ within a multiaction is replaced by an action c

Example of mCRL2 specification

Dining philosophers

eqn $K = 2$;

map K : **Pos**;

act $get, put, up, down, lock, free$: **Pos#Pos**;

eat : **Pos**;

proc

$Phil(n: \mathbf{Pos}) = get(n, n) . get(n, if (n == K, 1, n+1)). eat(n). put(n, n)$
 $. put(n, if (n == K, 1, n+1)) . Phil(n)$;

$Fork(n: \mathbf{Pos}) = sum m: Pos . up(m, n) . down(m, n) . Fork(n)$;

init $allow$ ({ $lock, free, eat$ },

$comm$ ({ $get/up \rightarrow lock, put/down \rightarrow free$ },



$Phil(1) || \dots || Phil(K) || Fork(1) || \dots || Fork(K)$);

Reo to mCRL2 (Constraint automata semantics)

- Data flow observed at a channel end = action
- Synchronous channel, synchronous drain
 - $Sync = A/B.Sync;$
- Non-deterministic synchronous lossy channel
 - $LossySync = (A/B + A).LossySync;$
- Asynchronous drain
 - $AsyncDrain = (A + B).AsyncDrain;$
- FIFO1
 - $FIFO1 = A.B.FIFO1;$
 - $FullFIFO1 = B.FIFO1;$
 - Alternative encoding: $FIFO1(f: Bool) = (\neg f \rightarrow A \diamond B).FIFO1(\neg f);$
- Replication node
 - $Replicator = X|Y|Z.Replicator;$
- Merge node
 - $Merger = (X|Z + Y|Z).Merger;$

Reo to mCRL2: Data support

act $A, B: \text{Data}$

- Synchronous channel
 - $\text{Sync} = \sum_{d \in \text{Data}} A(d)|B(d) . \text{Sync};$
- Synchronous drain
 - $\text{SyncDrain} = \sum_{d_1, d_2 \in \text{Data}} A(d_1)|B(d_2) . \text{SyncDrain};$
- Synchronous lossy channel
 - $\text{LossySync} = \sum_{d \in \text{Data}} (A(d)|B(d) + A(d)) . \text{LossySync};$
- Asynchronous drain
 - $\text{AsyncDrain} = \sum_{d \in \text{Data}} (A(d) + B(d)) . \text{AsyncDrain};$
- **Filter** 
 - $\text{Filter} = \text{sum} \sum_{d \in \text{Data}} (\text{exp}(d) \rightarrow A(d)|B(d) \diamond A(d)) . \text{Filter},$ where $\text{exp}(d)$ is a boolean expression
- **Transformer** 
 - $\text{Transformer} = \sum_{d \in \text{Data}} A(d)|B(\text{exp}(d)) . \text{Transformer};$
- Replication node
 - $\text{Replicator} = \sum_{d \in \text{Data}} X(d)|Y(d)|Z(d) . \text{Replicator};$
- Merge node
 - $\text{Merger} = \sum_{d \in \text{Data}} (X(d)|Z(d) + Y(d)|Z(d)) . \text{Merger};$

Reo to mCRL2: Data and time support

- FIFO1
 - $DataFIFO1 = struct\ empty?isEmpty\ |\ full(e:Data)?isFull;$
 - $FIFO1(f: DataFIFO1) = \sum_{d \in Data} isEmpty(f) \rightarrow A(d).Fifo1(full(d)) \diamond B(e(f)).FIFO1(empty))$
- T-timer with off- and reset- options
 - Reacts differently to different data inputs:
 - $DataTimer = struct\ reset?isReset\ |\ off?isOff\ |\ timeout\ |\ other(e: Data)?isOther$
 - Has two states
 - $State = struct\ OFF?isOFF\ |\ ON?isON$
 - State s (timer ON or OFF), current time x , timer delay t
 - $Timer(s: State, x: Real, t: Real) =$
 $isOFF(s) \rightarrow \sum_{d \in DataTimer} isOther(d) \rightarrow A(d).Timer(ON, 0, t) +$
 $isON(s) \rightarrow ((x < t) \rightarrow \sum_{d \in DataTimer}$
 $isReset(d) \rightarrow A(d).Timer(ON, 0, t) +$
 $isOff(d) \rightarrow A(d).Timer(OFF, x, t) +$
 $tick^c x.Timer(ON, x + 1, t))$
 $\diamond B(timeout).Timer(OFF, x, t)$

Reo to mCRL2: Global data types

- A connector should deal with any data items consumed by its source nodes
- Given a set of elementary data types DT_1, \dots, DT_n (e.g., inferred from web service interface specifications), the global data type is described as follows:
 - $Data = struct D_1(e_1: DT_1) | \dots | D_n(e_1: DT_n)$
- *Join node*
 - $Join = \sum_{d_1, d_2 \in Data} (X(d_1) | Y(d_2) | Z(tuple(d_1, d_2))).Join;$
- For m -join node $tuple(e_1: Data, e_2: Data, \dots, e_m: Data)$ is added to the Data description, e.g.,
 - $Data = struct D_1(e_1: DT_1) | \dots | D_n(e_1: DT_n) | tuple(e_1: Data, e_2: Data)$
- **Note:** expressions for filter and transformer channels become dependent on the structure of the Reo connector

Reo to mCRL2: Composition

Synchronize and hide actions corresponding to the connected channels




■ Problem

- Reduce the size of the state space while building the LTS for the mCRL2 specification of a Reo connector

■ Idea

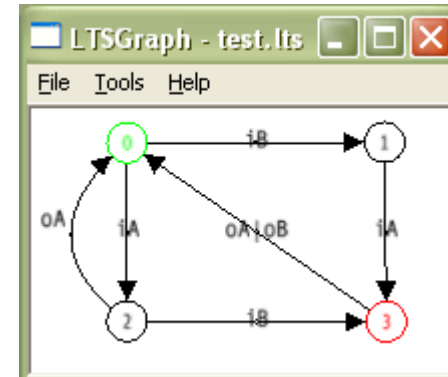
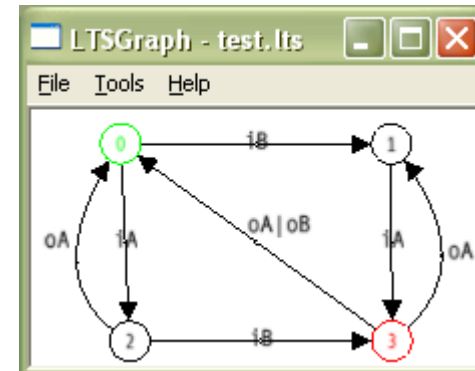
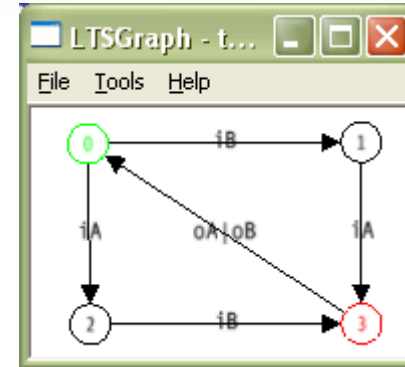
- Iterated connector construction



1. $P_0 = \partial_{\text{ends of connected channels}} (\Gamma_{\text{handshaking at Node}_1} (\text{Node}_1 \parallel \text{Sync}_1 \parallel \text{LossySync}_1 \parallel \text{LossySync}_2 \parallel \text{SyncDrain}_1))$
2. $P_1 = \partial_{\text{ends of connected channels}} (\Gamma_{\text{handshaking at Node}_2} (\text{Node}_2 \parallel \text{Sync}_2 \parallel \text{Sync}_3 \parallel P_0))$
3. $P_2 = \partial_{\text{ends of connected channels}} (\Gamma_{\text{handshaking at Node}_3} (\text{Node}_3 \parallel \text{Sync}_4 \parallel P_1))$
4. $P_3 = \partial_{\text{ends of connected channels}} (\Gamma_{\text{handshaking at Node}_4} (\text{Node}_4 \parallel \text{Sync}_5 \parallel P_2))$

Reo to mCRL2 (Intentional automata semantics)

- Two types of actions:
 - A channel end is ready to write/read (iA)
 - Data flow observed at a channel end (oA)
- Synchronous channel, synchronous drain
 - $Sync(a: Bool, b: Bool) =$
 $\neg a \rightarrow iA.Sync(\neg a, b) +$
 $\neg b \rightarrow iB.Sync(a, \neg b) +$
 $(a \wedge b) \rightarrow oA/oB.Sync(\neg a, \neg b);$
- Context independent lossy channel
 - $LossySync(a: Bool, b: Bool) =$
 $\neg a \rightarrow iA.LossySync(\neg a, b) +$
 $\neg b \rightarrow iB.LossySync(a, \neg b) +$
 $a \rightarrow (b \rightarrow oA/oB.LossySync(\neg a, \neg b) +$
 $oA.LossySync(\neg a, b));$
- Context dependent lossy channel
 - $LossySync(a: Bool, b: Bool) =$
 $\neg a \rightarrow iA.LossySync(\neg a, b) +$
 $\neg b \rightarrow iB.LossySync(a, \neg b) +$
 $a \rightarrow (b \rightarrow oA/oB.LossySync(\neg a, \neg b) +$
 $\neg b \rightarrow oA.LossySync(\neg a, b));$



Tool support

The screenshot displays the Eclipse SDK interface for a Reo diagram named 'loanRequest'. The diagram is a stateful process with the following components and transitions:

- Components:** Two 'Writer' components (each with 'requests=-1') and one 'Reader' component (with 'requests=-1').
- Transitions:**
 - start** (initial state) transitions to **checkClientProfileIn** and **login**.
 - login** transitions to **authorized** (condition: $e1(d) == \text{Alice}$).
 - checkClientProfileIn** transitions to **checkClientProfileOut** (condition: $\text{salary}(e1(d)) > 3 * \text{amount}(e1(d)) / (\text{period}(e1(d)) * 12)$).
 - checkClientProfileOut** transitions to **denied** (condition: $\text{salary}(e1(d)) < 2 * \text{amount}(e1(d)) / (\text{period}(e1(d)) * 12)$).
 - denied** transitions to **stop**.
 - processRequestIn** transitions to **processRequestOut** (condition: $\text{isSalarySufficient}$).
 - processRequestOut** transitions to **approved**.
 - approved** transitions to **stop**.

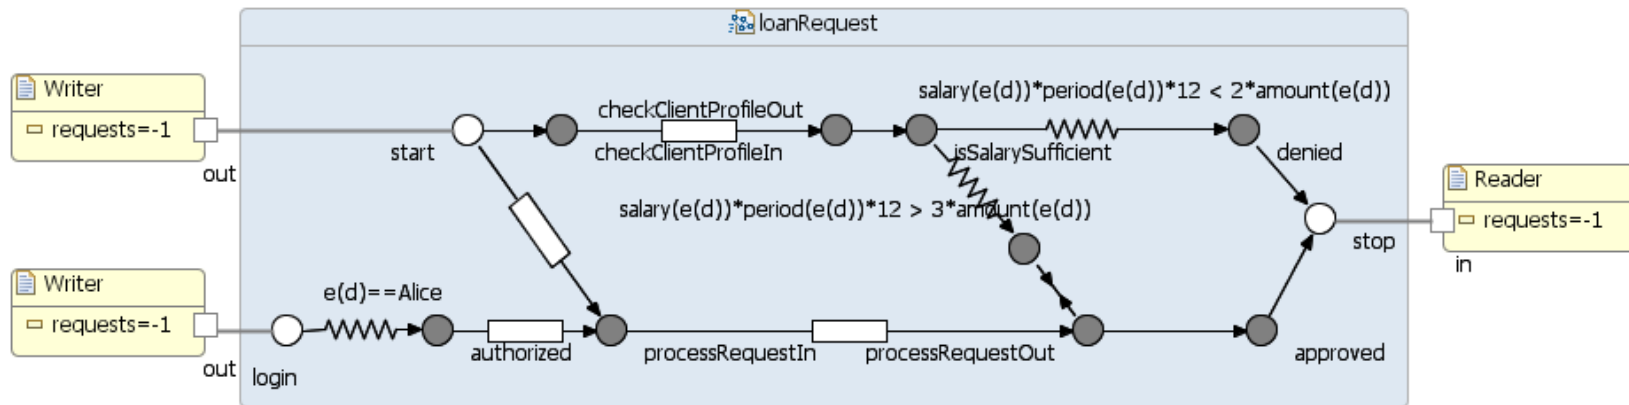
The bottom panel shows the mCRL2 encoding for the diagram:

```

mCRL2 Encoding:
Datatype:
Formula: [true*]<true>true
Options:  Network scope  With data  Intensional encoding  User-defined tau
Specification:
sort
Data = struct d1(e1 : DataWriter1)?isDataWriter1 | d2(e1 : DataWriter3)?isDataWriter3 | Tuple2(t1 : Data, t2 : Data)?isTuple2;
DataWriter1 = struct request(amount: Pos, salary: Pos, period: Pos)?isRequest;
DataWriter3 = struct Alice|Bob;
DataFIFO = struct empty | full(e : Data);

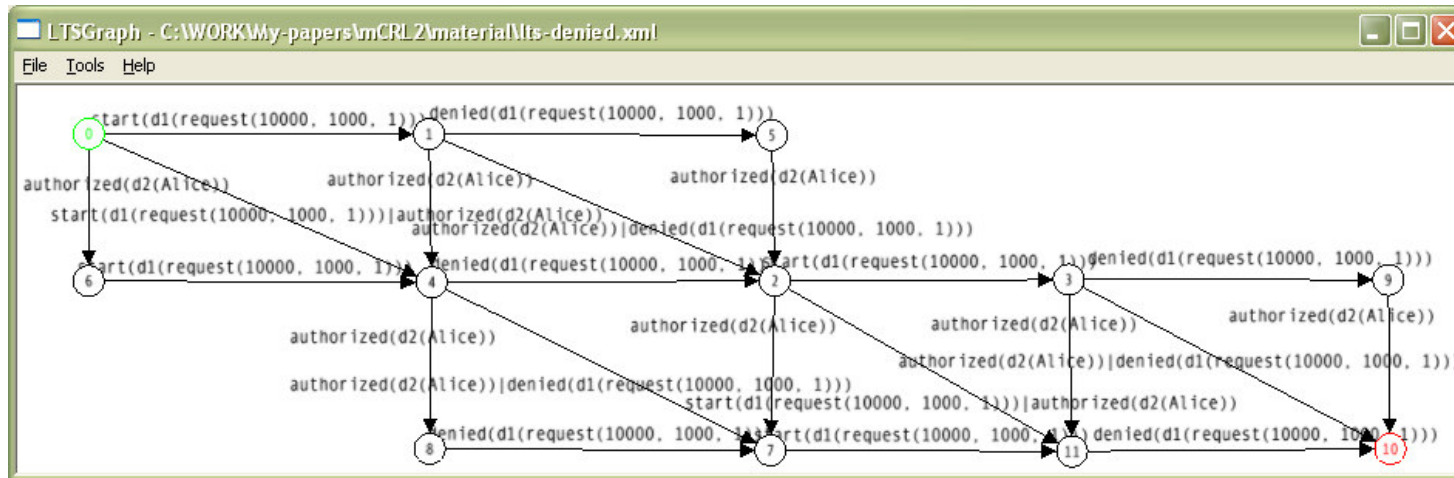
act
A', A'', Approved', Approved'', Approved0', Approved0'', Authorized', Authorized'', Authorized0', Authorized0'', B',
B'', CheckClientProfileIn', CheckClientProfileIn'', CheckClientProfileIn0', CheckClientProfileIn0'', CheckClientProfileOut', CheckClientProfileOut'', CheckClientProfileOut0', C
Denied0', IsSalarySufficient', IsSalarySufficient'', IsSalarySufficient0', IsSalarySufficient0'', IsSalarySufficient01', IsSalarySufficient01'', Login', Login'', Login0', Login0'', F
ProcessRequestIn', ProcessRequestIn0', ProcessRequestIn01', ProcessRequestIn01'', ProcessRequestIn011', ProcessRequestIn011'', ProcessRequestOut', ProcessRequestOut'', ProcessReques
Start', Start0', Start0'', Start01', Start01'', Stop', Stop'', Stop0', Stop0'', Stop01', Stop01'', login, approved, denied, authorized,
start, stop : Data;
    
```

Data-flow analysis with mCRL2

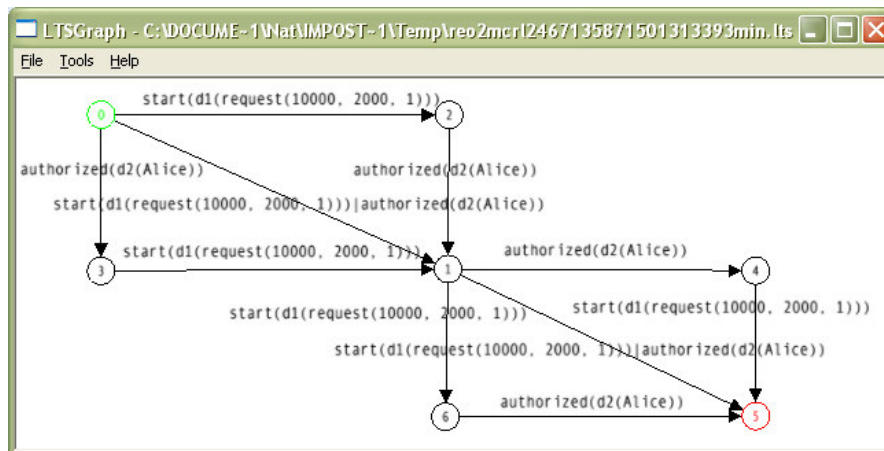


- μ calculus (property specification format for mCRL2)
 - No deadlock:
 - $[true^*]\langle true \rangle true$
 - No livelock:
 - $[true^*]\mu X.[\tau]X$
 - Loan request is always denied or approved:
 - Data agnostic:
 - $\nu X.[true]X \ \&\& \ \mu Y.([!(denied \ || \ approved)]Y \ \&\& \ \langle true \rangle true)$
 - Data aware:
 - $\nu X.[true]X \ \&\& \ \text{forall } d:\text{Data}. \ \mu Y.([!(denied(d) \ || \ approved(d))]Y \ \&\& \ \langle true \rangle true)$

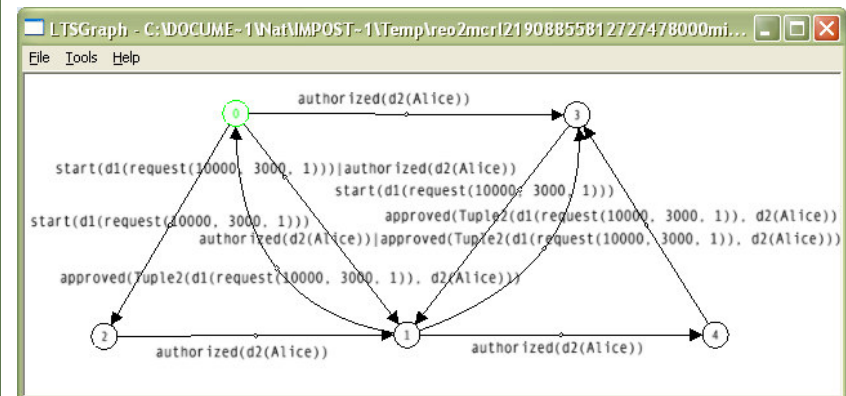
Loan request scenario: different instances



Salary = 1000



Salary = 2000



Salary = 3000

Comparison of model checking tools for Reo

- Vereefy (University of Dresden)
 - **Advantages:**
 - Developed for Reo and Constraint Automata
 - Counterexamples
 - **Disadvantages:**
 - No support for abstract data types
 - Global domain for all components
 - Primitive data constraint specification language (for filter channels)
- mCRL2 toolset
 - **Advantages:**
 - Powerful support for data
 - Very rich property specification format (μ calculus)
 - **Disadvantages:**
 - Hard to extract counterexamples
 - For infinite domains model checker often does not terminate (problems with algorithms for formulae rewriting)
 - Inability to define some useful data domains
 - Intentional model: state explosion problem
- CADP toolset (INRIA)
 - **Advantages:**
 - Fully compatible with mCRL2
 - Many useful tools (e.g., for performance evaluation)
 - **Disadvantages:**
 - License

Conclusions

- Full featured model checking of Reo connectors
 - Control + data flow analysis
 - Abstract data types
 - Complete tool support
 - Automated generation of mCRL2 code from graphical models
 - First implemented plug-in that deals with
 - Filter and transformer channels
 - Intentional semantics
 - Time channels
- Relation between Reo and ACP
- *Eclipse Coordination Tools* are an interesting toolset for business process and service composition analysis
 - Good alternative to Petri nets
 - Better fits service-oriented computing paradigm
 - *N. Kokash, C. Krause, E. de Vink: "Data-aware Design and Verification of Service Compositions with Reo and mCRL2", ACM SAC 2010, Technical track on Service-Oriented Architectures and Programming (SOAP)*

Future Work

- Work on tools usability
 - Analysis across several connectors
 - Obtaining data types from WSDL specifications
 - Other verification tools (e.g., CADP, simulators)
- Timed Reo
 - What properties can be checked with mCRL2
- Application to COMPAS case studies
 - Express compliance requirements and common BP properties in μ -calculus
 - **Problem:** property specification language is very expressive, but hard to use