

# Verifying System-Wide Properties of Industrial Component-Based Software

Thomas Neele, Marijn Rol, and Jan Friso Groote

Eindhoven University of Technology, The Netherlands  
{T.S.Neele, J.F.Groote}@tue.nl

**Abstract.** Analytical Software Design (ASD) enables model-based development of component software systems. Until now, functional verification of ASD systems is only possible on a per-component basis. There is no functional verification engine for ASD itself, so this verification relies on a translation of individual components to mCRL2, a process-algebraic model checker. We show how to extend the ASD-mCRL2 translation to support multiple components in order to enable checking of system wide functional properties. With our extended translation, we perform a case-study on a newly developed industrial system consisting of 26 communicating components. The results indicate that it is feasible to model check functional properties on this scale.

## 1 Introduction

Modern high-tech industry relies more and more on software to implement supervisory control logic. With the large number of software components in a typical machine, the software can become very complex. The industry not only wants software that meets high quality standards to assure safety and reliability, but the reduction of the costs and time of development also plays an important role. This is assured by model based software design accompanied with formal analysis where software problems are eradicated as soon as possible in the development process. Comparative research shows that it is possible to reduce the number of bugs by a factor 10 and the development time by a factor 3 [13].

Analytical Software Design (ASD) [2,3] is one of the model based engineering tools being used in industrial environments. Using ASD, software engineers develop models which can be checked for various properties such as deadlock/livelock freedom and interface compliance with a single press of a button. From the models, ASD generates executable code, e.g. Java or C++, that can be run in a production environment.

In contrast to many other tools that apply model checking techniques, ASD does not suffer severely from the *state-space explosion* problem. This is achieved through the application of compositional verification techniques: each component in a system is checked individually by comparing its implementation and interface using *failure divergence refinement* (FDR) [16]. A pleasant property of FDR is that deadlock/livelock freedom of each component guarantees deadlock/livelock freedom of the complete system. The ASD approach has been used

to develop systems with over 200 components [11] (more than 300 models if interface and design models are counted separately), where total verification takes around 20 minutes.

It is also possible to check a broader range of properties on single components through a translation to mCRL2 [4,10]. The process algebraic description language mCRL2 comes with a toolset for simulating, visualising, manipulating and model checking behavioural specifications. We call the functional properties that span a single component *local properties*.

There are however many global properties, also called ‘end-to-end’ properties, that are not covered by only checking local properties. Typical examples are:

- If the software control is instructed to manufacture a product, then the appropriate associated low level instructions are always issued.
- If one of the actuators reports an error, the control system always reports the error to the higher software layers.
- If the control software reports that the machine is off, it will never instruct any of its controlled actuators to move.

In this paper we report on how we verify such global properties on newly developed industrial control software. For this purpose we extend the existing mCRL2 translation to support multiple components. Firstly, communication has to be restricted to take place between the right components. Furthermore, important functionality that was implemented in C++, instead of ASD, needs to be translated manually to mCRL2. Finally, we must add several mechanisms to preserve the single-threaded execution as defined by the semantics, *i.e.*, we must ensure that only one component is active at a time.

We evaluate the approach by translating an ASD system that is newly developed and which consists of 26 components (together containing 5054 so called rule cases). On the resulting mCRL2 model, we check a complete set of end-to-end properties. Because the state-space of the system only consisted of 178 million states, we were able to establish whether each requirement was satisfied.

Our expectation was that we would encounter many hardships in the verification, especially because the state spaces would be excessively large. But the contrary turned out to be true. The state space remained well within acceptable limits for the available computer equipment. The reason for this appears to lie in the run-to-completion semantics employed in ASD, together with the strict use of interfaces. These coincide with design rules for systems to avoid the state space explosion [9]. There were a number of minor issues that had to be overcome, such as speeding up writing intermediate results to disk. It also turned out that applying a branching bisimulation reduction to the intermediate state space before applying model checking was much more time efficient than following the ordinary workflow.

Our conclusion is that it is very well possible to verify actual industrial-size software while it is under development. But for success, it needs to be written in an appropriate domain specific language whose semantics avoids a state space explosion.

*Overview* Section 2 introduces the basic concepts on which this paper is built. We explain our approach to multi-component translation in Section 3. The case study is introduced in Section 4 together with the properties we verify and Section 5 contains the results of the experiments we conducted. In Section 6, we give an overview of some related work. Finally, Section 7 presents a conclusion.

## 2 Background information

This section provides a short introduction to ASD, mCRL2 and the modal  $\mu$ -calculus, which form the bases of our approach.

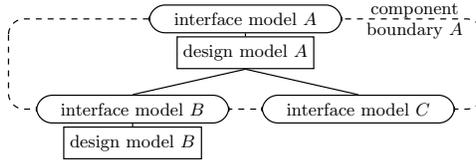
### 2.1 Analytical Software Design

Analytical Software Design (ASD) [3], developed by Verum, enables the development of software based on communicating components. The components are designed and verified using the ASD:Suite. Furthermore, the ASD:Suite can generate the executable code that can be used in a production environment. In ASD, there are two types of components:

- *Standard ASD component.* A standard component consists of an *interface* and a *design* model. The interface model specifies the externally visible behaviour of a component. It provides a more abstract view on the behaviour of a component. The design model specifies the inner working of a component, including how it interacts with lower level components. The design model always refines the interface model under failure-divergence refinement [16].
- *Foreign component.* A foreign component consists of only an interface model. It typically models the behaviour of a hardware component or another system that is implemented outside of ASD, *e.g.*, in C++.

If a component  $A$  relies on another component  $B$  for certain tasks, then we say  $A$  is a *client* of  $B$  and  $B$  is a *server* of  $A$ . Intuitively, an interface model serves as a contract on the behaviour of the corresponding design model. The interface model specifies exactly in which order a client can send calls to a server and which responses it can expect. Most of the decision logic of a component is contained in the design model. In ASD, the components must be structured as a tree, *i.e.*, a component cannot have more than one client. See Figure 1.

All ASD models are described using a formalism similar to *extended finite state machines*, which are state machines augmented with data. Software engineers develop these state-machines in the ASD:Suite in a format called *Sequence-based Specification* (SBS). A model can contain *state variables*, which store information about the state of the component. State variables can be booleans or any other finite enumeration type. Transitions can be guarded with expressions over the state variables, and state variables can be updated after every transition. A guarded transition can only be taken when the guard evaluates to *true* given the current values of the state variables. A row in a SBS is called a *rule case*.



**Fig. 1.** An example system composed of two standard components, *A* and *B*, and a foreign component *C*.

*Example 1.* We consider a component *A* that can be activated and deactivated and also paused and resumed. Component *A* is a client of component *B*, which always needs to be deactivated before component *A* can become inactive. A sequence-based specification of the design model of component *A* is given in Figure 2: it shows the four rule cases of state *Active* and one rule case of the state *Inactive*. When one of the actions *Pause* or *Resume* is performed, an empty reply (*VoidReply*) is returned and the variable *Enabled* is updated (rule cases 4 and 5). Component *A* can only be deactivated when it is not *Enabled* (rule cases 1 and 2). Upon deactivation, it first sends a message to component *B*, and only then deactivates itself by going to the state *Inactive*. Activation from the state *Inactive* happens in a similar way. We assume the actions not shown for one of the states are *blocked* in that state, *i.e.*, they cannot happen.  $\square$

Within ASD, we distinguish four types of communication: *call events*, *reply events*, *notification events* and *modelling events*. A call event happens when a client wants to request certain information or a certain action from one of its servers. While the server is handling that call event, it may choose to send one or more notification events to its client. These notifications are stored in the *notification queue* of the client until control is given back to the client. The notification queue will be explained in detail later. When the server is finished with a call event, it will always send a reply event to the client. When a reply contains data, we call it a *valued reply*, otherwise it is called a *void reply*.

| Case                   | Event        | Guard            | Actions  | Variable update | Target state |
|------------------------|--------------|------------------|--|-----------------|--------------|
| <b>State: Active</b>   |              |                  |  |                 |              |
| 1                      | A.Deactivate | enabled == false | B.Deactivate;<br>A.Deactivated;<br>A.VoidReply | -               | Inactive     |
| 2                      | A.Deactivate | Otherwise        | Illegal  | -               | -            |
| 3                      | A.Activate   | -                | NoOp   | -               | Active       |
| 4                      | A.Pause      | -                | A.VoidReply                                    | enabled = false | Active       |
| 5                      | A.Resume     | -                | A.VoidReply                                    | enabled = true  | Active       |
| <b>State: Inactive</b> |              |                  |  |                 |              |
| 6                      | A.Activate   | -                | B.Activate;<br>A.Activated;<br>A.VoidReply     | -               | Active       |
| 7                      | A.Deactivate | -                | NoOp   | -               | Inactive     |

**Fig. 2.** Example of a SBS for the design model of a component called *A*.

A modelling event typically represents low level input, such as an interrupt, which can be guaranteed to come, or which can incidentally take place. It can only be performed by a foreign component. Similar to notification events, they are sent from a server to a client. However, whereas notification events always happen as the result of a call event, modelling events only happen spontaneously.

There are two possible semantics for ASD: the *multi-threaded execution model* and the *single-threaded execution model* [11]. In this paper, we only consider the latter. In the single-threaded execution model events cannot happen in parallel, but only in sequence. Therefore, these semantics should define clearly in which order events are processed. According to the documentation of ASD, the following actions take place in order when a component receives an event from a client:

- All actions from the SBS rule case are processed in order.
- State variables are updated.
- The transition to the target state is taken.
- The notifications in the queue are processed. No events other than those caused by these notifications may occur before the queue has been emptied.
- A void or valued reply takes place to give control back to the client.

These rules are also referred to as *run-to-completion semantics*, meaning that a component completes all of its tasks before relinquishing control to another component.

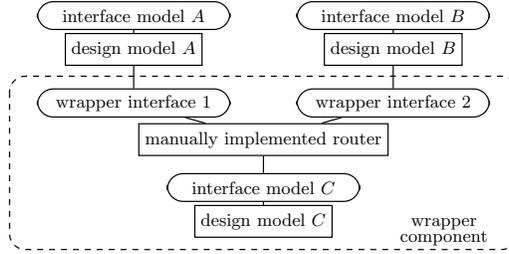
**Wrapper Components.** The tree structure of ASD is somewhat restrictive in practice, since it does not allow a component to have more than one client. That makes it impossible to implement bidirectional communication channels that have two clients, or database-like components that have many clients. One can work around this issue by manually implementing foreign components, as we explain below.

In the case study that we consider, wrapper components are used to implement symmetric communication channels, see Figure 3. In this case, there are three components  $A$ ,  $B$  and  $C$ . The components  $A$  and  $B$  are both a client of  $C$ , albeit through two separate interfaces. The router – written manually in C++ – implements both these interfaces. Since ASD is not aware of the connection between both interfaces, this structure does not violate ASD’s single-client constraint.

The router forwards all requests from components  $A$  and  $B$  to  $C$ , and also sends responses from  $C$  back to the correct client. For all requests component  $C$  receives from component  $A$ , it sends a notification to  $B$ , and vice versa. In this way, the wrapper component  $C$  serves as a bidirectional communication channel.

## 2.2 mCRL2

The language mCRL2 is a process-algebraic language [10] which can be analysed using the accompanying toolset [4]. The main aim of mCRL2 is model checking



**Fig. 3.** The structure of a wrapper component.

of parallel processes. Additionally, mCRL2 can generate, reduce, compare and visualise state-spaces.

The mCRL2 modelling language is very flexible and despite a limited set of language primitives, very expressive. Therefore, it is very well suited as a target language for automatic translations.

*Example 2.* To illustrate some of the concepts behind mCRL2, we consider the following specification.

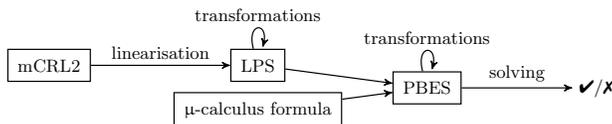
```

act tick, reset, press;
proc Clock( $n:\text{Nat}$ ) = tick.Clock( $n+1$ ) + reset.Clock(0);
      Button = press.reset.Button;
init allow({tick, reset, press},
           comm({reset|reset  $\rightarrow$  reset},
                Clock(0) || Button));

```

In this system, we have two processes, a *Clock* and a *Button*. The clock can perform an action *tick*, after which it increases the time (stored in parameter  $n$ ). When the button is pressed, it subsequently communicates with the clock via the action *reset*, and the clock resets its counter. Communication is enforced through the combination of the **allow** and **comm** operators, which in this case express that both *reset* actions must happen synchronously. For a more complete overview of the mCRL2 language, see [10].  $\square$

The typical mCRL2 workflow is depicted in Figure 4. We start out with a specification in the mCRL2 language. Such a specification can contain arbitrarily nested structures, therefore it is first linearised to a normal form, called a *linear process specification* (LPS). The standardised format of LPSs allows the application of transformations that may speed-up further analysis. An LPS can be combined with a functional property, specified in the modal  $\mu$ -calculus, to obtain a *parameterised Boolean equation system* (PBES). Again, we may choose to apply one or more transformations on this PBES. Solving the PBES will give us the answer to the model checking problem that we are considering. The solution of the PBES is true exactly if the investigated property holds for the specification.



**Fig. 4.** The mCRL2 model checking workflow

### 2.3 Modal $\mu$ -calculus

To express formal properties, the mCRL2 toolset relies on the modal  $\mu$ -calculus [12] with data, which is suitable to express virtually any conceivable behavioural property. It is far more expressive than LTL/CTL, but it is equally efficient when it comes to establishing those properties [5]. Here, we provide the core grammar of the  $\mu$ -calculus, *i.e.*, without the use of data:

$$\phi ::= \text{false} \mid \text{true} \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \phi \Rightarrow \phi \mid \langle a \rangle \phi \mid [a]\phi \mid \mu X. \phi \mid \nu X. \phi \mid X$$

Here  $a$  is an action,  $X$  is a fixpoint variable representing a set of states and  $\mu X. \phi$  and  $\nu X. \phi$  are the least and greatest fixpoints over  $X$  respectively. Formulae in the  $\mu$ -calculus are interpreted over labelled transition systems. The semantics is roughly as follows. The boolean operators have their usual semantics. The diamond modality  $\langle a \rangle \phi$  is *true* if and only if an  $a$ -step is possible after which  $\phi$  holds. The box modality  $[a]\phi$  expresses that after every possible  $a$ -step,  $\phi$  must hold. The least fixpoint  $\mu X. \phi$  is true for the smallest set of states  $X$  such that  $\phi$  holds for all states in  $X$ . Note that  $X$  can occur in  $\phi$ . Dually,  $\nu X. \phi$  is true for the largest set  $X$  that satisfies  $\phi$ . The least fixpoint operator expresses that a property must be valid within a finite number of iterations, whereas the greatest fixpoint also allows for infinitely repeating behaviour.

The mCRL2 toolset also allows specifying modalities with sets of actions via so-called *action formulas*: *true* represents the set of all actions and *false* represents the empty set of actions. Supported operators are union, intersection and inverse ( $\bar{a}$  is the set of all actions other than  $a$ ). Furthermore, sequences of actions can be specified with *regular formulas*, which give the possibility to concatenate sequences (with the  $.$  operator), take their union (operator  $+$ ) or iterate over them (operator  $*$ ). The action formula  $a^*$  represents zero or more occurrences of the sequence  $a$ , and  $\text{true}^*$  represents any sequence of actions. For example,  $[a.b]\text{false}$  means that a sequence  $ab$  is not possible (as if it is possible, *false* must hold in the resulting state) and  $\langle a^*.b \rangle \text{true}$  means that a sequence consisting of zero or more  $a$ 's followed by a  $b$  is possible. Action formulas and regular formulas can always be expressed using the fixed point operators, but they are generally more convenient to specify concrete behavioural properties.

In this work, we do not consider  $\mu$ -calculus formulae with data. In mCRL2, data variables can be bound in quantifiers, *i.e.*,  $\exists$  or  $\forall$ , or as parameter of a fixpoint variable. They are used in conditions, and as parameters of actions and fixpoint variables. See [10] for a complete overview of the modal  $\mu$ -calculus.

### 3 Approach

The existing translation from ASD to mCRL2 is only capable of translating the models within one component boundary at a time [11]. This translation yields two mCRL2 specifications: one containing the topmost interface model and one containing the design model and the interface models below it. On the latter specification, we can check local properties that concern only that component.

We defined a new translation that yields a single mCRL2 specification that represents the behaviour of the complete system. The new translation takes as input all the design models and also the interface models of foreign components. Before we introduce the challenges introduced by the new translation in detail, we first introduce the basic single-component translation.

#### 3.1 Translating single components

Due to the expressivity of mCRL2, ASD components can be mapped almost directly to mCRL2. First, for every state of the component, a recursive process is created. This process carries one parameter for each state variable of the ASD model. For each rule case, this process has one *action summand*, which contains the condition, actions, variable updates and target state deduced from the rule case. Furthermore, the mCRL2 specification contains a *Queue* process that represents the behaviour of the notification queue. A formal definition of the translation can be found in [11].

*Example 3.* We revisit the component *A* from Example 1. For this model, the corresponding mCRL2 specification looks as follows:

$$\begin{aligned}
A\_Active(enabled:Bool) = & \\
& (enabled \approx false) \rightarrow A\_deactivate\_called \cdot invoke\_B\_Deactivate \cdot \\
& \quad emit\_A\_Deactivated \cdot sendReply(void) \cdot A\_Inactive(enabled) + \\
& (enabled \approx true) \rightarrow A\_deactivate\_called \cdot Illegal \cdot \delta + \\
& A\_activate\_called \cdot NoOp \cdot A\_Active(enabled) + \\
& A\_Pause\_called \cdot sendReply(void) \cdot A\_Active(false) + \\
& A\_Resume\_called \cdot sendReply(void) \cdot A\_Active(true) \\
A\_Inactive(enabled:Bool) = & \\
& A\_activate\_called \cdot invoke\_B\_Activate \cdot emit\_A\_Activated \cdot \\
& \quad sendReply(void) \cdot A\_Active(enabled) + \\
& A\_deactivate\_called \cdot NoOp \cdot A\_Inactive(enabled)
\end{aligned}$$

For both states in the SBS, we have a process definition (*A\_Active* and *A\_Inactive* respectively), which carry the state variable *enabled*. The different types of events each have a prefix or suffix to distinguish them: received call events have the suffix *called*, sent call events have the prefix *invoke* and notification events have the prefix *emit*. After an illegal event, the process deadlocks (operator  $\delta$ ).  $\square$

### 3.2 Communication

In the existing single-component translation, a void reply is represented with the action *sendReply(void)*. Since the scope of this translation is very limited – it only translates one design model and several interface models from one component boundary (cf. Figure 1) – synchronization on this action can only take place in one way: between the design model and one of the interface models. However, in the multi-component setting, we have to explicitly enforce synchronization to happen between the proper components. Therefore, every occurrence of an action *sendReply* does not only have an argument for the type of the reply and zero or more arguments to carry data, but also two arguments to indicate the source and destination of the reply. In this way, only those components will synchronize on that action. The same approach is applied to notifications.

### 3.3 Manual translations

The automatic translation to mCRL2 cannot handle wrapper components, since the router is implemented in C++ instead of ASD. The behaviour of a wrapper component as a symmetric communication channel is essential to the behaviour of the complete system, therefore, it is desirable that the wrapper components are also present in the translation of the complete system. We manually translated the router to mCRL2, because this is far more efficient than performing an automatic translation from C++.

There is another component, called *AsyncCall*, of which the behaviour is manually translated from C++ to mCRL2. The component *AsyncCall* can be requested by any component to send a response at some later time. This is a workaround such that ASD components can awaken themselves to finish residual duties. Internally, these requests are stored in a queue in *AsyncCall*. The queue can only contain one request per component and components also have the option to cancel their request.

We remark that the wrapper components are partially generic and the *AsyncCall* component is completely generic. Therefore, they do not need to be implemented from scratch when verifying several systems. Ideally, the mCRL2 specification of the wrapper components can be generated by the same program that generates their C++ implementation.

### 3.4 Queues

As defined in the semantics of ASD, every component contains its own queue to store notifications. This implies that the complete mCRL2 specification will have a queue for every design model in the system. To ensure that the run-to-completion semantics is preserved, we add a lock to every queue. A queue is unlocked exactly when the corresponding component is active processing a call from a client or a modelling event. In this way a client can only process the content of a queue when it is active.

### 3.5 Framework

Not only the ASD components themselves, but also the outside world, which we will refer to as the *framework*, should behave according to certain constraints. For example, the framework cannot send out another call or modelling event when the system did not yet finish the previous task. To encode this, we add the following features to our translation:

- A new action *emptyQ* that can only be executed when all queues are empty; all queues synchronise on this action.
- An additional process *Thread* controls the sending of calls to the uppermost components and sending of modelling events by the foreign components. At the moment it sends a call or modelling event, it checks whether the queues are empty. It can only send another message after the previous one completed processing.

## 4 Case study

We perform a case study to investigate the feasibility of our approach and to determine its applicability to industrial-size systems. The case study is based on a real-life ASD system found in the model stack of our industrial partner. Components and events have been renamed for confidentiality reasons. Figure 5 shows the high-level structure of the system. The system consists of two loosely-coupled subsystems, called *A* and *B*. *A* and *B* have to cooperate to execute an action together, which we will call *Exec*. The clients of *A* and *B* independently decide whether they are ready to do so. Moreover, after a client has requested for *Exec* to be performed, it can repeal its decision by sending a *Cancel* message. Both sides synchronize using wrapper components located on the *cancel layer* and the *control layer*. The cancel layer consists of the components responsible for cancelling the execution of *Exec* and the control layer is responsible for performing *Exec*. The cancel and control layers, including their direct server components, consist of 14 components. The complete system consists of 26 components, which contain 5054 rule cases in total.

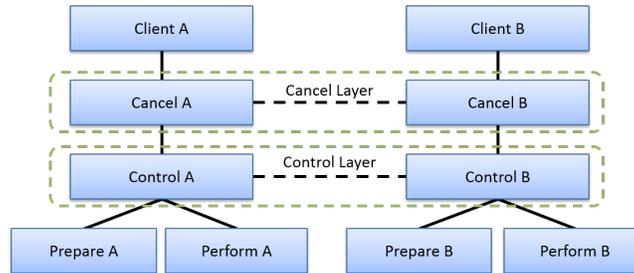
In the initial state, the clients of *A* and *B* can decide whether they want to perform *Exec* or cancel it. When both clients request to do *Exec* without sending a *Cancel* at some point, then this action is performed. If at least one client requests a cancellation before they both request *Exec*, then the action *Cancel* is performed. If a client asks for *Exec* to be cancelled after *Exec* has started, then it will not be cancelled. Both clients will receive a notification when an *Exec* or a *Cancel* has been completed. After a client requests *Exec*, it will immediately start preparing for *Exec* to happen. When *Exec* or *Cancel* has finished, the system should return to the initial state. We call the process between the first message and the performing of *Exec* or *Cancel* a *round*.

## 4.1 Subsystems

In our analysis, we consider four different variants of this system, to get a rough idea of the scalability. Firstly, we have the *layers subsystem*, which consists of the cancel layer, the control layer, the components directly below them and the wrapper components in between. We consider two variants of the layers subsystem: the regular implementation and an implementation where no *Cancel* request can be performed. Secondly, we have two variants of the complete system: one that does not allow errors to occur and one that does allow errors. We will call the former *good-weather behaviour* (GWB) and the latter *bad-weather behaviour* (BWB). The bad-weather behaviour system is rather rudimental, which means that clients can raise errors which are subsequently dealt with by the system to cause the least disturbance in the normal process operation.

## 4.2 Properties

After consulting the domain experts, we identified several system-wide properties that the system under study should adhere to. For the layers subsystem, we have eight properties, which are listed in Table 1. Furthermore, we have three properties that involve the complete system; they are listed in Table 2. Note that some properties are symmetric for both clients; in that case we only listed one of the two  $\mu$ -calculus formulae. Many properties are concerned with behaviour from the moment that neither client has sent a message yet until the moment that *Exec* or *Cancel* is performed. Since we are dealing with a system that runs continuously, we should not only check what happens from the initial state, but in every round. Therefore, we formulated a property in the  $\mu$ -calculus that expresses whether the system is at the start of a round (property ‘Initial state’ in Tables 1 and 2). This formula is used within other properties to check behaviour in all rounds. Most action names in these properties should be self-explanatory. The actions *outwardReply* and *outwardNotification* respectively represent a reply and a notification sent to one of the two clients. The inserting of a notification into the queue of a certain component is represented by the action *raiseNotification*.



**Fig. 5.** The structure of the case study system. The dashed lines indicates bidirectional communication through wrapper components.

**Table 1.** Properties of the layers system written in modal  $\mu$ -calculus

| Property   | Formula  |
|--|--|
| 0 Initial state  | $(\langle A\_Request\_Exec \rangle true \wedge \langle A\_Request\_Cancel \rangle true \wedge \langle B\_Request\_Exec \rangle true \wedge \langle B\_Request\_Cancel \rangle true)$                                   |
| 1 Cannot do <i>Exec</i> without two requests                 | $[true^*](Initial\ state \Rightarrow [(A\_Request\_Exec^* + B\_Request\_Exec^*).Exec]false)$   |
| 2 Cannot request execution after a cancel request            | $[true^*.A\_Request\_Cancel.(outwardNotification(Cancelled))^*.A\_Request\_Exec]false$   |
| 3 Must perform execution after two execute requests          | $[true^*](Initial\ state \Rightarrow [A\_Request\_Exec.(A\_Request\_Cancel + B\_Request\_Cancel)^*.B\_Request\_Exec]\mu X.([\overline{Exec}]X \wedge \langle true \rangle true))$                                      |
| 4 Raise cancel notification after two cancel requests        | $[true^*.A\_Request\_Cancel.(outwardNotification(Cancelled))^*.B\_Request\_Cancel]\mu X.([\overline{outwardNotification(Cancelled)}]X \wedge \langle true \rangle true)$   |
| 5 After a cancel request, <i>Cancel</i> is performed         | $[true^*](Initial\ state \Rightarrow [A\_Request\_Cancel + B\_Request\_Cancel]\mu X.([\overline{A\_Cancel}]X \wedge \langle true \rangle true))$   |
| 6 Cannot make multiple <i>Exec</i> requests in a round       | $[true^*.A\_Request\_Exec.(outwardNotification(Exec.finished))^*.A\_Request\_Exec]false$   |
| 7 No synchronization error during an execution               | $[true^*](Initial\ state \Rightarrow [A\_Request\_Exec.(B\_Request\_Exec)^*.B\_Request\_Exec.(A\_Request\_Exec + B\_Request\_Exec)^*.Sync\_Error(A)]false)$  |
| 8 Synchronization error after requesting an execute too soon | $[true^*](\langle A\_Request\_Exec \rangle true \wedge \langle B\_Get\_Results \rangle true) \Rightarrow [(B\_Get\_Results)^*.A\_Request\_Exec]\mu X.([\overline{Sync\_Error(A)}]X \wedge \langle true \rangle true))$ |

In these properties, we use several common patterns. First, it is very common to write a property of the shape  $[true^*]\varphi$ , meaning that after any sequence of actions,  $\varphi$  has to hold. Building on that, the property  $[true^*.a]false$  expresses that action  $a$  cannot occur anywhere and the property  $[true^*.a.b^*.c]false$  expresses that after every action  $a$ , an action  $b$  must happen before the action  $c$  happens.

A more complex, but important, pattern is  $\mu X.([\overline{a}]X \wedge \langle true \rangle true)$ , which means that action  $a$  unavoidably happens within a finite amount of steps. The intuition is as follows: as long as we do something other than  $a$  (subformula  $[\overline{a}]$ ), we recurse through variable  $X$ . That may only happen finitely often, due to the least fixpoint ( $\mu X$ ). Therefore, we must at some point end up in a state where actions other than  $a$  are not possible. This state cannot be a deadlock, since that is explicitly forbidden by  $\langle true \rangle true$ . Hence, we end up in a state where only  $a$  actions can be done, and at least one  $a$  is possible. This is the same as saying that ultimately,  $a$  must be done.

## 5 Results

Our approach to verifying the properties differs from the standard workflow presented in Section 2.2. In that workflow, at least part of the state space is explored while solving the PBES. That means that we would need to perform state-space exploration several times, once for every property. Instead, we generate the state-space directly from the LPS (cf. Figure 4), and we subsequently minimise the state-space under divergence-preserving branching bisimulation us-

**Table 2.** Properties of the full system written in modal  $\mu$ -calculus

| Property  | Formula   |
|---|---|
| 0 Initial state   | $\langle (A\_Request\_Exec.(\overline{Protocol\_Error(A)})^*.outwardReply(VoidReply))true \wedge$ $\langle A\_Request\_Cancel.(\overline{Protocol\_Error(A)})^*.outwardReply(VoidReply) \rangle true \wedge$ $\langle B\_Request\_Exec.(\overline{Protocol\_Error(B)})^*.outwardReply(VoidReply) \rangle true \wedge$ $\langle B\_Request\_Cancel.(\overline{Protocol\_Error(B)})^*.outwardReply(VoidReply) \rangle true$ |
| 1 Must perform execution after two execute requests               | $[true^*(Initial\ state \Rightarrow [A\_Request\_Exec.(\overline{A\_Request\_Cancel} + \overline{B\_Request\_Cancel})^* .$ $B\_Request\_Exec]\mu X.([\overline{Exec}]X \wedge \langle true \rangle true))$  |
| 2 Prepare steps are done before perform steps                     | $[true^*(Initial\ state \Rightarrow [A\_Request\_Exec.(\overline{A\_Request\_Cancel} + \overline{B\_Request\_Cancel})^* .$ $B\_Request\_Exec.(\overline{raiseNotification(A\_Prepare\_Step\_Done)})^* .Exec]false)$   |
| 3 Perform steps are done before raising an execution notification | $[true^*(Initial\ state \Rightarrow [A\_Request\_Exec.(\overline{A\_Request\_Cancel} + \overline{B\_Request\_Cancel})^* .$ $B\_Request\_Exec.(\overline{raiseNotification(A\_Perform\_Step\_Done)})^* .outwardNotification(Exec\_finished)]false)$  |

ing the Groote-Jansen-Keiren-Wijs algorithm [8]. From the state-space we construct a Boolean equation system, which can be solved to obtain an answer to the model checking question.

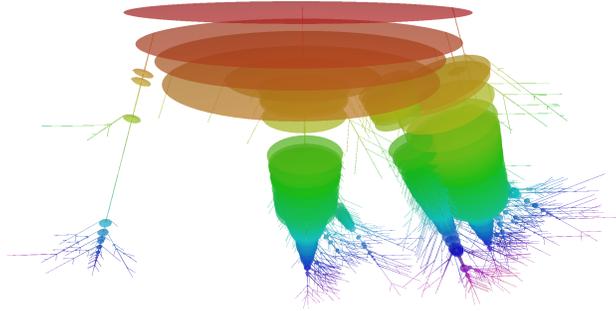
To run the experiments, we used a machine with multiple Xeon E5520 processors (56 cores in total), clocked at 2.27GHz and 935GB of memory. The mCRL2 version we installed has Git commit hash 73241e378e<sup>1</sup>. The mCRL2 analysis tools are all single threaded.

Table 3 gives an overview of the time required for state-space generation and the size of each of the state spaces. The time reported does not include the time required for bisimulation reduction, which is about 45 minutes for the BWB system. The full system under bad-weather behaviour is almost on the limit of what can be generated in a reasonable amount of time with 178.6 million states. At the same time, bisimulation reduction is very effective, and manages to bring the number of states back to 12.4 million.

<sup>1</sup> Sources are available at <https://github.com/mCRL2org/mCRL2>

**Table 3.** Time required for state-space generation and number of states and transitions for the systems before and after divergence-preserving branching bisimulation minimisation.

| System                  | time (s) | mem   | #states     | #transitions | After bisim. red. |              |      |
|-------------------------|----------|-------|-------------|--------------|-------------------|--------------|------|
|                         |          |       |             |              | #states           | #transitions | %red |
| Layers no <i>Cancel</i> | 7        | 28MB  | 9,107       | 9,472        | 7,085             | 7,422        | 22.2 |
| Layers subsystem        | 46       | 63MB  | 109,608     | 114,310      | 55,361            | 58,338       | 49.5 |
| Complete system GWB     | 14,713   | 1.7GB | 17,179,798  | 19,098,495   | 3,787,974         | 4,298,103    | 75.0 |
| Complete system BWB     | 154,397  | 14GB  | 178,603,107 | 196,784,882  | 12,451,325        | 14,879,416   | 93.0 |



**Fig. 6.** Visual representation of the minimised state space of the layers subsystem.

Figure 6 shows a visualisation of the minimised state space of the layers subsystem. The initial state is at the top and every disk represents a (non-deterministic) choice. Initially, the system contains little branching behaviour (the red, yellow and green parts at the top). Only deeper in the state space, there is more choice to perform different actions (blue and purple parts at the bottom). The low amount of branching can be ascribed to the run-to-completion semantics of ASD.

Table 4 records for each of the four variants of the system the average time required to check one of the properties. All properties hold, except when checked on the bad-weather behaviour system, since that has not been fully implemented. The time required for the full system is significant: almost one hour for the good weather version and close to five hours per property for the bad weather system.

While running these large experiments, we observed that a lot of time and memory is spent on storing and loading intermediate files from disk. For a typical property of the BWB system, `lts2pbes` spends more than three quarters of the time on storing the PBES on disk. This problem could be by-passed by implementing an integrated tool that combines the functionality of `lts2pbes` and `pbessolve`. The amount of memory required to verify a property of the BWB system is roughly 180GB; this peak is also reached while writing the output in `lts2pbes`.

**Table 4.** Average time spent to verify a single property on each of the four (sub)systems.

|                        | time (s)         |               |                 |                 |
|------------------------|------------------|---------------|-----------------|-----------------|
|                        | Layers no cancel | Layers System | Full System GWB | Full System BWB |
| <code>lts2pbes</code>  | 2.06             | 9.82          | 2,317           | 14,126          |
| <code>pbessolve</code> | 0.44             | 3.74          | 890             | 4,150           |
| Total                  | 2.50             | 13.56         | 3,207           | 18,276          |

## 6 Related work

The successor of ASD is Dezyne<sup>2</sup>, also developed by Verum. Similar to the ASD-mCRL2 translation, there is also a translation from Dezyne to mCRL2 [17]. This translation is also limited to single components, so it does not support verification of end-to-end properties.

mCRL2 has also been used in other studies to analyse systems with a very large state space. For example, in [1], the train control system ERTMS Hybrid Level 3 is analysed with mCRL2. They apply the same workflow as we do: minimise the transition system with bisimulation before checking any property. The largest state space they verified contains close to 34 million states.

Remenska *et al.* [14] also work with an automated translation. They convert the behaviour captured in UML2.0 sequence diagrams to mCRL2 specifications. Their technique is applied on DIRAC, the computing grid framework of CERN's LHCb experiment. Although the state-space is too large to generate completely, they do find a counter-example to the desired property with depth-first search.

The idea of generating code from formal models that have been checked with model checking is also applied in [6]. The authors present a tool called DLC, Distributed LNT Compiler, which can produce C code from an LNT specification. The generated code is suited for running on multiple machines concurrently, synchronization between the machines is achieved with a rendezvous protocol. LNT specifications can be analysed with the existing tools from the CADP toolset [7].

## 7 Conclusion

We showed how to translate a component system implemented in ASD to mCRL2. This enables checking of end-to-end properties on ASD, which is important for mission-critical software. Furthermore, we demonstrated with a case-study that this approach is applicable to an actual industrial system. For the two variants of the layers subsystem, the time and space required to run state-space generation and verification is sufficiently small to enable interactive verification during development. This is due to the semantics of ASD that avoids a major state-space explosion. The results exceeded our own expectations and give us hope that model checking of complete systems can be applied more often in industrial settings.

One of the challenges that needs to be tackled before wider adoption of this approach is possible, is the complexity of modal  $\mu$ -calculus. Currently, specifying properties with  $\mu$ -calculus requires expertise, and it is not uncommon for formulas to contain mistakes. A possible solution is to supply developers with natural-language templates in which they enter the correct action names. The corresponding formula will then be generated from the template [15].

---

<sup>2</sup> See <https://www.verum.com/>, accessed 30-08-2018

## References

1. M. Bartholomeus, B. Luttik, and T. A. C. Willemse. Modelling and Analysing ERTMS Hybrid Level 3 with the mCRL2 Toolset. In *FMICS 2018*, volume 11119 of *LNCS*, pages 98–114, 2018.
2. G. H. Broadfoot. ASD Case Notes: Costs and Benefits of Applying Formal Methods to Industrial Control Software. In *FM 2005*, volume 3582 of *LNCS*, pages 548–551, 2005.
3. G. H. Broadfoot and P. J. Hopcroft. Analytical software design. Technical report, Verum Consultants B.V., 2003.
4. S. Cranen, J. F. Groote, J. J. A. Keiren, F. P. M. Stappers, E. P. de Vink, W. Wesselink, and T. A. C. Willemse. An Overview of the mCRL2 Toolset and Its Recent Advances. In *TACAS 2013*, volume 7795 of *LNCS*, pages 199–213, 2013.
5. S. Cranen, J. F. Groote, and M. A. Reniers. A linear translation from CTL\* to the first-order modal  $\mu$ -calculus. *Theor. Comput. Sci.*, 412(28):3129–3139, 2011.
6. H. Evrard and F. Lang. Automatic distributed code generation from formal models of asynchronous processes interacting by multiway rendezvous. *Journal of Logical and Algebraic Methods in Programming*, 88:121–153, 2017.
7. H. Garavel, F. Lang, R. Mateescu, and W. Serwe. CADP 2011: A Toolbox for the Construction and Analysis of Distributed Processes. *STTT*, 15(2):89–107, 2013.
8. J. F. Groote, D. N. Jansen, J. J. A. Keiren, and A. J. Wijs. An  $O(m \log n)$  Algorithm for Computing Stuttering Equivalence and Branching Bisimulation. *ACM Trans. on Comput. Logic*, 18(2), 2017.
9. J. F. Groote, T. W. D. M. Kouters, and A. Osaiweran. Specification guidelines to avoid the state space explosion problem. *Softw. Test., Verif. Reliab.*, 25(1):4–33, 2015.
10. J. F. Groote and M. R. Mousavi. *Modeling and Analysis of Communicating Systems*. MIT Press, 2014.
11. R. Jonk. The semantics of ALIAS defined in mCRL2. Master’s thesis, Eindhoven University of Technology, 2016.
12. D. Kozen. Results on the propositional  $\mu$ -calculus. *Theoretical Computer Science*, 27(3):333–354, 1982.
13. A. Osaiweran, M. Schuts, J. Hooman, J. F. Groote, and B. J. van Rijnsoever. Evaluating the effect of a lightweight formal technique in industry. *STTT*, 18(1):93–108, 2016.
14. D. Remenska, J. Templon, T. A. C. Willemse, P. Homburg, K. Verstoep, A. Casajus, and H. Bal. From UML to process algebra and back: An automated approach to model-checking software design artifacts of concurrent systems. In *NFM 2013*, volume 7871 of *LNCS*, pages 244–260, 2013.
15. D. Remenska, T. A. C. Willemse, J. Templon, K. Verstoep, and H. E. Bal. Property specification made easy: Harnessing the power of model checking in UML designs. In *FORTE 2014*, volume 8461 of *LNCS*, pages 17–32, 2014.
16. A. W. Roscoe. On the expressive power of CSP refinement. *Formal Aspects of Computing*, 17(2):93–112, 2005.
17. R. van Beusekom, J. F. Groote, P. Hoogendijk, R. Howe, W. Wesselink, R. Wieringa, and T. A. C. Willemse. Formalising the Dezyne Modelling Language in mCRL2. In *FMICS/AVoCS 2017*, volume 10471 of *LNCS*, pages 217–233. Springer, 2017.