A Case Study in Parallel Verification of Component-Based Systems

N. Beneš, I. Černá, J. Sochor, P. Vařeková and B. Zimmerova[1,2]

Faculty of Informatics, Masaryk University
Brno, Czech Republic

Abstract

In large component-based systems, the applicability of formal verification techniques to check interaction correctness among components is becoming challenging due to the concurrency of a large number of components. In our approach, we employ parallel LTL-like model checking to handle the size of the model. We present the results of the actual application of the technique to the verification of a complex model of a real system created within the *CoCoME Modelling Contest* [18]. In this case study, we check the validity of the model and the correctness of the system via checking various temporal properties. We concentrate on the component-specific properties, like local deadlocks of components, and correctness of given use-case scenarios.

Keywords: Component-based systems, formal verification, parallel model checking.

1 Introduction

During the last decade, software industry has seriously started to take advantage of component-based software development as an alternative to existing software development techniques. Component-based development proposes to assemble software systems from reusable components, possibly in a hierarchical manner. This helps to significantly reduce development costs, but brings the issue of correctness of such systems, especially if components are delivered by different vendors.

In this paper, we present a practical application of parallel verification to a large component-based system designed within the *CoCoME Modelling Contest* [15]. In the contest, a number of teams were asked to create a detailed model of a common component-based system to make their modelling approaches comparable. While in [18], we present our model of the CoCoME system, this paper complements the work by verifying the model. In verification, we concentrate on properties of the final model like correctness of given use-case scenarios, local deadlocks of components,

---

and response properties. Besides these we discuss how the verification helped us to check the validity of the model during modelling.

As a modelling language for component-based systems we use *Component-Interaction automata* (or *CI automata* for short) [6,8] which allow very precise and detailed description of communication among system components. System properties are specified in an extended version of the action-based linear time logic LTL, called *CI-LTL*. For the verification itself we use the automata-based model checking algorithms implemented in the parallel model checking tool *DiVinE* [4,10]. We advocate the choice of a parallel tool by a tremendous size of the model given by concurrency of components in the system.

A short description of the CoCoME Modelling Contest is given in Section 2 followed by an outline of the CI automata modelling language and the CI-LTL logic in Section 3. Section 4 introduces the model we have created within the contest, and Section 5 lists required properties and use-case scenarios including their verification. Finally, Section 6 discusses the results and experience gained during the verification.

2 CoCoME Modelling Contest

In order to leverage component-based system design to build correct and dependable component-based systems, researchers have developed various formal and semi-formal component models which concentrate on different yet related aspects of component modelling [13,7,5,12,2,11]. The main goal of the *CoCoME (Common Component Modelling Example) Modelling Contest* [15] was to evaluate and compare the practical application of existing component modelling approaches and techniques on a common modelling example, which was designed to comprise a large number of various aspects and modelling issues that can be identified in different types of component-based systems.

The modelling example, called *Trading System*, serves to handle sales in a chain of supermarkets. Its functionality includes the interaction with the cashier at the cash desk, like product scanning, price lookup, cash/card payment, and bill printing, as well as accounting the sale at the inventory, or determining whether an express cash desk is needed in the store. Furthermore, the Trading System deals with ordering goods from wholesalers, and generating various kinds of reports. The system is an open system, designed to interact with external components representing users of the system (cashiers and managers) and a bank application.

The Trading System was implemented as a Java application where components correspond to packages in the source code. The Java source code (125 Java classes in total) served as a detailed specification of the system for the modelling teams to prevent ambiguities in the interpretation of the corresponding high-level specification. The component structure of the application up to depth four is depicted in Figure 1. The figure includes an *id* number for each primitive component in the system. If a component is assigned more than one *id*, it consists of several sub-components with these *id*s. The *id*s are used as numerical names of the components in automata labels.
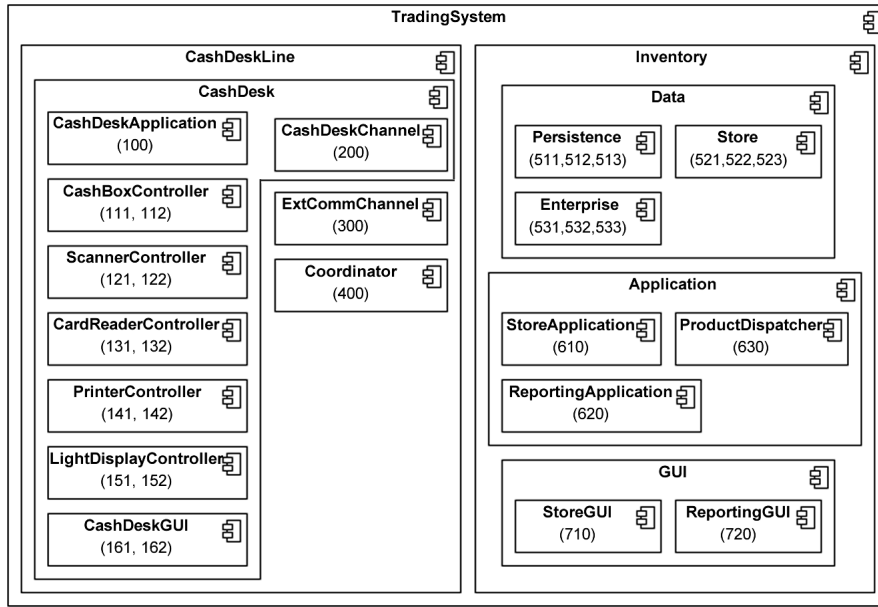
Fig. 1. CoCoME Trading System overview

3 Foundations

3.1 Modelling language

To model behaviour of component-based systems we use the *CI automata* language [6,8]. The language models each component as a labelled transition system with structured labels and a hierarchy of component names. The transition label articulates which components communicate on an action, and the hierarchy of names represents the architectural structure of the component.

A *CI automaton* is a 5-tuple $\mathcal{C} = (Q, Act, \delta, I, H)$ where Q is a finite set of states, *Act* is a finite set of *actions*, $\Sigma = ((S_H \cup \{-\}) \times Act \times (S_H \cup \{-\})) \setminus (\{-\} \times Act \times \{-\})$ is a set of *labels*, $\delta \subseteq Q \times \Sigma \times Q$ is a finite set of *labelled transitions*, $I \subseteq Q$ is a nonempty set of *initial states*, and H is a structured tuple representing a hierarchy of component names where the set of component names is denoted S_H.
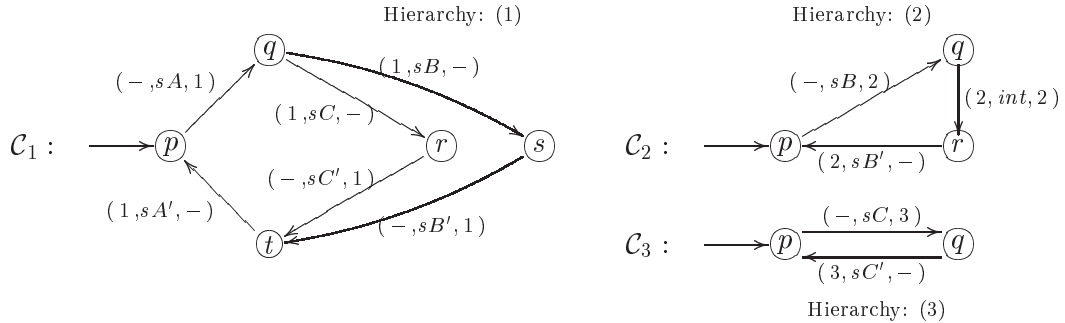

Fig. 2. Three examples of CI automata

The labels have semantics of input, output, or internal, based on their structure. In the triple, the middle item represents an action name, the first item represents a name of the component that outputs the action, and the third item represents

a name of the component that inputs the action. Examples of three CI automata are in Figure 2. Each of them represents a model of behaviour of a basic component. For example, $(-, sA, 1)$ in C_1 signifies that the component with numerical name 1 inputs an action sA (a request for a service $sA()$), and $(1, sA', -)$ in C_1 signifies that the component 1 outputs an action sA' (a response for the service $sA()$).

To compose components into a higher-level component a composition operator is defined. Automata can be composed together using a parametrizable composition operator $\otimes^{\mathcal{F}}$, which composes a given finite set of automata with respect to the set of *feasible labels* $\mathcal{F}$. Given a set of labels $\mathcal{F}$, the operator composes the set of CI automata into a product automaton allowing only those transitions from the product that have labels from $\mathcal{F}$. In the product, the components cooperate either by interleaving of their original transitions, or by simultaneous execution of two complementary transitions (with labels $(n_1, a, -)$, $(-, a, n_2)$) which results into a new internal transition (with label (n_1, a, n_2)). An example of a composite automaton is in Figure 3. A wider range of composition operators is defined in [6,8].
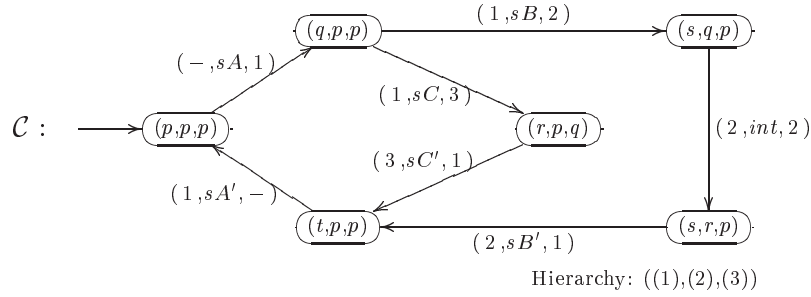

Fig. 3. A composite CI automaton $\mathcal{C} = \otimes^{\mathcal{F}}\{C_1, C_2, C_3\}$ where C_1, C_2, C_3 are in Fig. 2, and $\mathcal{F} = \{(-, sA, 1), (1, sA', -), (1, sB, 2), (2, sB', 1), (2, int, 2), (1, sC, 3), (3, sC', 1)\}$

3.2 Temporal logic

For property specification, we use a slightly modified version of the linear temporal logic LTL [14] which we refer to as *CI-LTL*. CI-LTL is designed to express properties about occurring component interaction (i.e. labels in automata), but also about possible component interaction (i.e. label enabledness).

Syntax. For a given set of labels, formulas of CI-LTL are defined as

(1) $\mathcal{P}(l)$ and $\mathcal{E}(l)$ are formulas, where l is a label.

(2) If Φ and Ψ are formulas, then also $\Phi \wedge \Psi$, $\neg \Phi$, $\mathcal{X} \Phi$ and $\Phi \mathcal{U} \Psi$ are formulas.

(3) Every formula can be obtained by a finite number of applications of steps (1) and (2).

Other operators can be defined as shortcuts: $\Phi \vee \Psi \equiv \neg (\neg \Phi \wedge \neg \Psi)$, $\Phi \Rightarrow \Psi \equiv \neg (\Phi \wedge \neg \Psi)$, $\mathcal{F} \Phi \equiv \text{true } \mathcal{U} \Phi$ (Future), $\mathcal{G} \Phi \equiv \neg \mathcal{F} \neg \Phi$ (Globally).

Semantics. Let $\mathcal{C} = (Q, Act, \delta, I, H)$ be a CI automaton. We define a *run* of $\mathcal{C}$ as an infinite sequence $\sigma = q_0, l_0, q_1, l_1, q_2, \ldots$, where $q_i \in Q$, and $\forall i . (q_i, l_i, q_{i+1}) \in \delta$. We further define:

- $\sigma(i) = q_i$ (i-th state of σ)

- $\sigma^i = q_i, l_i, q_{i+1}, l_{i+1}, q_{i+2}, \ldots$ (i-th sub-run of σ)
- $\mathcal{L}(\sigma, i) = l_i$ (i-th label of σ)

CI formulas are interpreted over runs and the satisfaction relation $\models$ is defined as

$$
\begin{aligned}
\sigma &\models \mathcal{E}(l) & &\Longleftrightarrow \exists q \,.\, \sigma(0) \xrightarrow{l} q \\
\sigma &\models \mathcal{P}(l) & &\Longleftrightarrow \mathcal{L}(\sigma, 0) = l \\
\sigma &\models \Phi \wedge \Psi & &\Longleftrightarrow \sigma \models \Phi \text{ and } \sigma \models \Psi \\
\sigma &\models \neg\,\Phi & &\Longleftrightarrow \sigma \not\models \Phi \\
\sigma &\models \mathcal{X}\,\Phi & &\Longleftrightarrow \sigma^1 \models \Phi \\
\sigma &\models \Phi \,\mathcal{U}\, \Psi & &\Longleftrightarrow \exists j \in \mathbb{N}_0 \,.\, \sigma^j \models \Psi \text{ and } \forall k \in \mathbb{N}_0, k < j \,.\, \sigma^k \models \Phi
\end{aligned}
$$

Informally, formula $\mathcal{E}(l)$ is true in all states of the system where the interaction represented by the label l can possibly happen. Formula $\mathcal{P}(l)$ is true for a run whenever the interaction represented by the label l is actually happening as the very first transition of the run.

3.3 Model checking and verification tool

For model checking CI-LTL properties, the automata-based algorithm [17] is slightly modified in the way a formula is translated into a Büchi automaton. Automaton has a special alphabet formed by doubles (set of labels, label). The items correspond to the two operators $\mathcal{E}(l)$ and $\mathcal{P}(l)$. Apart from that, the model checking algorithm remains the same as in the case of standard LTL (accepting cycle detection) and therefore it has the same complexity.

The tool DiVinE, which we use for the verification, provides several LTL model checking algorithms. In our case study, the algorithm OWCTY is employed. This algorithm always generates the whole reachable state space of the model and its time complexity is higher than that of simple reachability if it is the case that an accepting cycle is found. However, it was chosen because of its efficiency in distributed setting [3].

The verifications presented in this paper have been performed on a cluster of ten 2.60 GHz Intel Pentium 4 Linux workstations with 3800 MB of RAM, interconnected with a 100Mbps Ethernet and using the Message Passing Interface (MPI) library. The chosen number of computers is explained and justified in Section 6.

4 Model of the Trading System

Within the *CoCoME Modelling Contest* [15], we have created a detailed model of the Trading System in terms of component interaction using CI automata [18]. The model in a textual notation is available at [16]. The model consists of 140 primitive automata (59 in the CashDeskLine part, and 81 in the Inventory part), composed hierarchically into 34 composite automata up to 6 levels of depth. The Trading System model is complemented by several models of cashiers and managers, who interact with the system, and specify various usage profiles under which properties

of the system are checked.[3] Each usage profile/scenario (all provided within the contest) corresponds to a correct behaviour of a user operating the system.

We have experimented with all usage scenarios. However, for the clarity of the presentation, we employ only one usage profile underlying the properties studied in the paper. It is the scenario describing one sale assisted by a cashier. This scenario represents the most complex usage profile described in [15], and it is connected to a large number of component-specific properties that can be checked on the behaviour of the system that is implied by the scenario. In the scenario, the cashier first starts the sale, then scans items (in a loop), finishes the sale and receives the payment. It can select cash or card payment, where the cash payment is followed by entering the received amount and returning change, and the card payment with scanning the card and entering PIN.

Besides the users, the system interacts with a bank application to exchange information during card payments. We suppose that the bank can perform any correct scenario, i.e. it is anytime able to receive requests and for each request it returns a response. We simulate this by leaving the communication with the bank open.

State space of the model. As mentioned above, the Trading System model is composed out of 140 primitive automata hierarchically assembled into 34 composite automata. Even if the size (number of states) of individual primitive automata is moderate, the size of the complete state space is immense due to the concurrency in component behaviour. An attempt to generate the complete state space on a cluster of twenty computers finished unsuccessfully with 322 millions of states demanding for 60 GB of memory in total. The computation took 13 400 seconds. Although this might seem slow, it is not surprising because the state space generation involves computation of the successor states. Such computation is a complex task, which needs to take into account the hierarchical composition of components and the feasible labels that can be propagated up in the hierarchy and are part of the composite automata.

Even if the complete model is unfeasibly large, for the verification of the model under the given usage scenario, the model is composed with an automaton representing the user. This restricts possible behaviours and decreases the state space. The size of the model with the cashier mentioned earlier is 749 340 reachable states and 3 181 473 reachable transitions.

5 Verification of the model

In this section, we discuss some of the properties that were checked on the model, and present verification results. We concentrate on the properties that are specific to component-based systems and emerged from the requirements on the Trading System and discussions with other teams. A part of the contribution of this paper is the identification of such a set of properties defining correctness issues in component-based systems, their formalization in terms of temporal logics, and demonstration of the feasibility and efficiency of their automatic verification in parallel settings.

---

[3] Only if we know, for instance, that a sale proceeds correctly including the payment, it is meaningful to check that all purchased goods were correctly taken off in the inventory.

Moreover, in the CoCoME Modelling Contest, a number of requirements were specified in terms of use-case scenarios. Use-case scenarios define a behaviour of the system in response to a given usage profile. Verification of use-case scenarios is studied after the other properties in this section, and is followed by discussion on the importance of formal verification, to check the validity of the model during the modelling process. The section concludes with experimental results studying the effect of parallelization on the verification.

5.1 Basic properties

As the basic properties, we present two properties demonstrating the capability of the *CashDeskChannel* component in the Trading System to broadcast events to the components that subscribed for them.

Property 1 (Unwanted duplicity). When the *CashDeskChannel (200)* receives a request to broadcast the *SaleSuccessEvent* via $(100, publishSaleSuccess-Event, 200)$, the event is going to be delivered to all subscribers $(200, onEventSale-Success, X)$ at most once. In the property, as well as in the following properties, action names are shortened to the sequence of first letters of their sub-words, e.g. *publishSaleSuccessEvent* becomes *pSSE*.

(a) $\quad \mathcal{G}\left(\mathcal{P}(100, pSSE, 200) \Rightarrow \neg\left[\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ (\mathcal{P}(200, oESS, 142) \wedge \right.\right.$
$$\mathcal{X}\left[\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 142)]\right)]\right)$$

(b) $\quad \mathcal{G}\left(\mathcal{P}(100, pSSE, 200) \Rightarrow \neg\left[\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ (\mathcal{P}(200, oESS, 162) \wedge \right.\right.$
$$\mathcal{X}\left[\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 162)]\right)]\right)$$

property	states	transitions	memory	time	result
prop1a	749 340	3 181 473	533 MB	67 s	holds
prop1b	749 340	3 181 473	535 MB	67 s	holds

The data in the table refer to the model composed with the appropriate property automaton. The column **memory** represents the total memory needed by all workstations in verification of the property. Note that the number of states of the model composed with the property is, in this case, equal to the number of states of the original model. This interesting fact is explained in Section 6.

Property 2 (Guaranteed delivery). Whenever the *CashDeskChannel (200)* receives a request to broadcast the *SaleSuccessEvent*, the event is going to be delivered to all subscribers $(200, onEventSaleSuccess, X)$ at least once, or an exception occurs $(200, exceptionPublishSaleSuccessEvent, 100)$.

$$\mathcal{G}\left[\mathcal{P}(100, pSSE, 200) \Rightarrow ([BOTH \wedge \neg EXC] \vee [NONE \wedge EXC])\right]$$

where

$BOTH = [\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 142)] \wedge [\neg \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 162)]$

$$NONE = (\neg\, [\neg\, \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 142)]) \wedge (\neg\, [\neg\, \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ \mathcal{P}(200, oESS, 162)])$$
$$EXC = \neg\, \mathcal{P}(100, pSSE, 200)\ \mathcal{U}\ (200, ePSSE, 100)$$

property	states	transitions	memory	time	result
prop2	749 340	3 181 473	533 MB	69 s	holds

5.2 Local deadlocks of components

In component-based systems, many components coexist in parallel. Hence deadlock of some of them cannot be detected as halting of the whole system. We understand a *local deadlock* of a component as a state from which the component *cannot* move further. This situation requires the *enabledness* $\mathcal{E}$ operator, otherwise we could only express that it *does not* move further. The following two properties describe a local deadlock of a component on a particular service call, and the third property specifies a local deadlock with respect to any action.

Property 3 (Local deadlock on one action). It cannot happen that the *Store-Application (610)* is ready to call `getTransactionContext()` but never can do so because its counterpart *Persistence (511)* is never ready to accept the call.

$$[\mathcal{F}\ \mathcal{P}(610, gTC, -)]\ \vee\ \mathcal{G}\ [\mathcal{E}(610, gTC, -) \Rightarrow \mathcal{F}\ \mathcal{E}(610, gTC, 511)]$$

property	states	transitions	memory	time	result
prop3	778 100	3 298 237	539 MB	69 s	holds

This property helped us to evaluate one of our modelling decisions. As the service `getTransactionContext()` activates a new instance of the component *TransactionContextImpl*, where only a limited number of instances can be active at any time, this property allows us to check that the bound on the number of instances that are ready to be activated is sufficient.

Note that this property requires the presence of the $(610, getTransaction\text{-}Context, -)$ label, which symbolizes an attempt of the component 610 to ask for a new transaction context, in the model. However, this is an output label that, according to the specification, must synchronize with a complementary input label before the model is complete, and be restricted from the model. Therefore, for the purpose of verification of this property, we modify the model in a way that this label is not restricted by the composition. However, to keep the verification faithful, the property is defined in a way that the runs with the label $(610, getTransactionContext, -)$ on them are not traversed during verification (as they have only informative character). Hence the increase in the size of the model (via not omitting the label) influences neither the state-space traversed during verification, nor the fidelity of the model.

Property 4 (Local deadlock on one action). It cannot happen that the

CashDeskApplication (100) is ready to send a notification to the *CashDeskChannel (200)* saying that it received the *SaleStartedEvent*, but the *CashDeskChannel* is never ready to accept the notification.

$$[\mathcal{F}\ \mathcal{P}(100, oESS'', -)] \vee \mathcal{G}\ [\mathcal{E}(100, oESS'', -) \Rightarrow \mathcal{F}\ \mathcal{E}(100, oESS'', 200)]$$

property	states	transitions	memory	time	result
prop4	749 343	3 181 479	532 MB	67 s	holds

The *CashDeskChannel (200)* in the system is not allowed to accept notifications before it delivers events to all subscribers. If some of the subscribers would be constantly refusing to accept the event, it could block other components that already accepted the event and want to notify the channel. As the property is valid, this cannot happen in the system (on the *SaleStartedEvent*).

Property 5 (Local deadlock on any action). It cannot happen that the *Persistence (511)* for StoreApplication becomes deadlocked (cannot make any action).

$$\mathcal{G}\ \mathcal{F}\ (ENABLED_{511})$$

where $ENABLED_{511} = \mathcal{E}(610, gPC, 511) \vee \mathcal{E}(620, gPC, 511) \vee \ldots \vee \mathcal{E}(511, eIA, 620)$, that is a disjunction of formulas of type $\mathcal{E}(label)$ for all labels the *Persistence (511)* participates in.

property	states	transitions	memory	time	result
prop5	1 498 679	7 805 074	689 MB	563 s	does not hold

The violation of the property means that the system gets into a state from which the component is no more able to perform any computation. This can happen for three reasons: (1) it gets stuck in its internal computation, (2) the environment refuses to accept its calls, or (3) the environment does not wish the component to compute anything for it any more. After a further analysis of the model we learned that the last case is true because in the usage profile, we suppose that only one sale is accomplished. Hence the system is not supposed to execute forever. This property demonstrates that one needs to be careful when interpreting the result from the local deadlock verification. More, it allows the reader to see the memory and time needed to verify a property that does not hold.

5.3 Blocking of components

A stricter version of local deadlocks, which is very interesting in component-based settings, is temporary *blocking* of a component because of the non-readiness of its counterpart to accept its calls. This property is considered the core issue of correct-

ness of component-based systems in several component-based models (SOFA [1], Interface automata [9]).

Property 6. It cannot happen that the *StoreApplication (610)* wants to begin a transaction $(610, beginTransaction, -)$ calling the *TransactionContextImpl (511)*, which is not right in the current state ready to accept it.

$$[\mathcal{F}\,\mathcal{P}(610, bT, -)] \,\vee\, \mathcal{G}\,\neg\,[\mathcal{E}(610, bT, -) \wedge \neg\,\mathcal{E}(610, bT, 511)]$$

property	states	transitions	memory	time	result
prop6	749 340	3 181 473	532 MB	67 s	holds

Note that we require the existence of the $(610, beginTransaction, -)$ label in the model. For the purpose of this verification, we modify the model in a way similar to the case with property 3. Even here, the resulting state space does not change, due to the nature of the property.

Property 7. It cannot happen that the *CashDeskApplication (100)* is ready to send a notification to the *CashDeskChannel (200)* saying that it received the *SaleStartedEvent*, but the *CashDeskChannel* is not right in the current state ready to accept the notification.

$$[\mathcal{F}\,\mathcal{P}(100, oESS'', -)] \,\vee\, \mathcal{G}\,\neg\,[\mathcal{E}(100, oESS'', -) \wedge \neg\,\mathcal{E}(100, oESS'', 200)]$$

property	states	transitions	memory	time	result
prop7	1 498 671	6 362 935	688 MB	534 s	does not hold

The property is a more strict version of the property 4. While the property 4 shows that the *CashDeskChannel (200)* always sends all copies of the *SaleStartedEvent* and gets into the state where it is ready to start accepting notifications, this property shows that it may take a while before the channel gets ready. However, this is not an error in the system. It correctly reflects the nature of the channel.

5.4 Loop issues

In our model, many cycles/loops can be found. Each loop can complete a run that enters it but never exits. In software systems, however, most of the loops in models result from `for` or `while` cycles that are traversed only finitely many times. The problem is that the number of traversals of the for/while cycles in the system is usually not known in advance—it is computed at run-time. Hence the cycles need to be modelled as loops, which by nature have no bound on the number of traversals. This can cause non-realistic results of properties verification. The properties should be verified only on the runs that follow selected loops only finitely many times.

Property 8. Whenever the *ProductDispatcher (630)* call `queryStoreById()` on the *Store* for *ProductDispatcher (523)* via $(630, queryStoreById, 523)$, it gets a response $(523, queryStoreById', 630)$ at some point in the future.

$$\mathcal{G} \left[\mathcal{P}(630, qSBI, 523) \Rightarrow \mathcal{F} \, \mathcal{P}(523, qSBI', 630) \right]$$

property	states	transitions	memory	time	result
prop8	750 684	3 186 705	533 MB	200 s	does not hold

In the counterexample, one of the components gets into a loop (representing a `for` cycle with a finite but unknown number of iterations possible) that it never exits. Hence the counterexample represents a run that is not real in the system. However, as there is no natural way to remove the run from the model for the reasons above, we modify the property in a way that it misses this run, thus forcing verification of the original property only on *fair* runs.

Property 9. Whenever the *ProductDispatcher (630)* calls `queryStoreById()` on the *Store (523)* for *ProductDispatcher*, it gets a response at some point in the future, if the progress of the system is forced by transitions of the *Store (523)*, which cannot get into invalid infinite loop.

$$\mathcal{G} \left[(\mathcal{P}(630, qSBI, 523) \wedge \mathcal{G} \, \mathcal{F} \; MOVE_{523}) \Rightarrow \mathcal{F} \, \mathcal{P}(523, qSBI', 630) \right]$$

where $MOVE_{523} = \mathcal{P}(610, qLSI, 523) \vee \mathcal{P}(620, qASI, 523) \vee \ldots \vee \mathcal{P}(630, qSI, 523)$, that is a disjunction of formulas of type $\mathcal{P}(label)$ for all labels the *Store (523)* participates in.

property	states	transitions	memory	time	result
prop9	750 684	3 186 705	534 MB	67 s	holds

Note that although the state-space size of the model composed with property 8 is the same as that of property 9, the verification time is larger in the first case. This is due to the nature of the verification algorithm, as mentioned in Section 3.3.

5.5 Use-case scenarios

In the verification of use-case scenarios, we are given an assumption on the usage profile of the system, and we want to guarantee that a particular behaviour is present in the response of the system. A use-case scenario is defined as a sequence of interactions (labels). It can be either complete (all labels are listed) or partial (given labels can be interleaved with other labels). In component-based systems, where the searched behaviour can be interleaved by behaviour of independent components in the system, the partial scenarios are of higher interest. This section presents results of verification of the three most complex (partial) scenarios defined in [15].

In contrast to the other verified properties, the use-case scenarios do not state that for all paths, some property holds (as is usual in the LTL model checking), but they state that there is a path, along which some property holds (namely the property representing the sequence of labels). This can be verified with the same methods, just by negating the property. Note that the properties representing the use-case scenarios are so large (their descriptions were over 100 lines long) that we do not give their formal representation here. However, they are a part of the model, which is available at [16].

UC scenario 1. CashPayment The scenario reflects cooperation of system components to successfully accomplish purchase of goods finished with cash payment.

UC scenario 2. Unsuccessful CardPayment The scenario describes system reactions to a sale finished with card payment that is refused by the bank.

UC scenario 3. Successful CardPayment The scenario describes component interaction following a successful sale finished with card payment.

property	states	transitions	memory	time	result
uc1	19 362 460	81 959 821	4 204 MB	5 141 s	scenario found
uc2	11 670 924	49 165 124	2 694 MB	3 203 s	scenario found
uc3	11 680 736	49 202 320	2 698 MB	3 098 s	scenario found

5.6 Validity of the model

During modelling, we needed to abstract from aspects of the system that could make the size of the model unmanageable, while staying confident about the safety of the abstractions. Two types of abstractions were considered: simplification of the internal behaviour of primitive components, and simplification of the communicational scheme. Regarding the communication among components, we evaluated serialization of selected parallel service calls and changing of some asynchronous calls to synchronous. The serialization was considered both on required (calling services) and provided (serving calls) side. This significantly reduced the state space, while causing no harm when the service calls were independent and their ordering had no effect on further behaviour of the system. Verification helped us to evaluate a number of serialization and synchronisation decisions via checking the validity of the model after the modification.

When checking the validity of the model, we worked with a set of properties based mainly on the use-case scenarios and test cases defined in the CoCoME Modelling Contest. We also tested the model for deadlocks, because we experienced that violation of the model validity often results in deadlock situations, either global or local.

prop	1	2	5	10	15	19
prop1a	186 MB	225 MB	339 MB	533 MB	736 MB	890 MB
prop1b	187 MB	226 MB	339 MB	535 MB	734 MB	896 MB
prop2	187 MB	225 MB	340 MB	533 MB	729 MB	888 MB
prop3	192 MB	231 MB	345 MB	539 MB	736 MB	898 MB
prop4	187 MB	225 MB	341 MB	532 MB	730 MB	890 MB
prop5	341 MB	379 MB	494 MB	689 MB	889 MB	1 052 MB
prop6	187 MB	225 MB	339 MB	532 MB	730 MB	893 MB
prop7	341 MB	379 MB	495 MB	688 MB	885 MB	1 050 MB
prop8	187 MB	226 MB	341 MB	533 MB	730 MB	889 MB
prop9	186 MB	225 MB	339 MB	534 MB	729 MB	889 MB
uc1	—	3 844 MB	3 989 MB	4 204 MB	4 411 MB	4 582 MB
uc2	2 303 MB	2 356 MB	2 491 MB	2 694 MB	2 895 MB	3 057 MB
uc3	2 305 MB	2 358 MB	2 493 MB	2 698 MB	2 900 MB	3 059 MB

Table 1
Memory consumption depending on number of computers.

5.7 Effect of parallelization

This section presents two tables summarizing the time and space consumption of the verification of presented properties, depending on the number of computers used in the computation. This was done in order to analyse the effect of the parallelization of the model checking algorithm and to choose the best number of workstations for detailed experiments. Table 1 shows the total memory used by all workstations, while Table 2 presents the time needed for the verification. A dash "—" indicates that the computation did not finish successfully (ran out of available memory).

6 Experience and discussion

In this section, we share our modelling and verification experience, discussing some of the results and observations we have achieved.

Characteristics of the model. As the number of components in the Trading System is quite large, and our modelling language expresses component concurrency through interleaving, the model suffers from state space explosion. More, the size of the reachable state space does not grow evenly during the hierarchical composition of components, but it changes dramatically. The reason for the irregular changes of the state space is that a composite automaton does not need to be larger than the automata it is composed of. We have observed cases, where the number of reachable states has been dramatically reduced by the composition. This is due to

prop	1	2	5	10	15	19
prop1a	586 s	312 s	130 s	67 s	48 s	41 s
prop1b	584 s	309 s	129 s	67 s	47 s	42 s
prop2	575 s	311 s	130 s	69 s	48 s	41 s
prop3	595 s	324 s	135 s	69 s	49 s	45 s
prop4	576 s	312 s	129 s	67 s	49 s	41 s
prop5	4 731 s	2 556 s	1 087 s	563 s	375 s	311 s
prop6	577 s	311 s	129 s	67 s	48 s	41 s
prop7	4 690 s	2 526 s	1 051 s	534 s	366 s	311 s
prop8	1 732 s	930 s	387 s	200 s	141 s	114 s
prop9	570 s	311 s	129 s	67 s	49 s	39 s
uc1	—	24 581 s	10 035 s	5 141 s	3 624 s	2 896 s
uc2	27 507 s	14 619 s	6 159 s	3 203 s	2 103 s	1 721 s
uc3	27 164 s	14 578 s	6 239 s	3 098 s	2 190 s	1 743 s

Table 2
Time consumption depending on number of computers.

the parametrized operator that can delimit possible behaviour in the composition. This fact can complicate the estimation of the number of states for a given model. But on the other hand, it can be exploited to produce a smaller model out of a large one, as was demonstrated in this case study, where the large Trading System model has been restricted by the usage profile of the cashier.

Deadlocks in the model. After deciding on the model for verification, in the validation phase, we have checked the model for global and local deadlocks. We have learned that the existence of deadlock states often signals a modelling error. A few *global* deadlocks were found. By careful investigation, we found that these deadlocks correspond to a behaviour reflecting that two components decide to receive messages from one of the event channels in an incorrect order, thus blocking each other. As we were not provided with the specification of the event channels, we can treat this finding in two ways. Either the deadlock reveals an error in the system, or it reflects an unrealistic behaviour, i.e. the system guards that the components receive messages in the right order. We decided to treat the runs leading to the deadlock states as unrealistic, and ignore them during verification. This is done implicitly in our verification method, because it verifies *infinite* runs only.

Local deadlocks and component-blocking properties. Interesting observations were made in verifying the *local* deadlocks and their more strict form, the component-blocking properties. We have verified many pairs of such properties and have found a strong relation between the two kinds. Mostly, it was either the case

that both properties were satisfied, or none of them was. The reasons are similar to those explained after property 5, that is, the environment does not wish the components to compute anything any more. We have, though, found a few cases, when the *local deadlock* property holds, but the *blocking* property does not, and we have presented one of them. Note that both kinds of properties take advantage of the *enabledness* $\mathcal{E}$ operator without which they could not have been formulated.

Size of the model/property composition. As may be noticed in Section 5, in some of the presented cases the state-space size remains (nearly) the same when the model is composed with the property automaton. This interesting fact deserves an explanation. The property automata are generated with the effort to make the resulting composition as small as possible. Then in case of some properties (such as safety and request/response properties), for every state of the model in the composition, there is a unique state in the property automaton. Hence the composition with the property does not influence the size of the model.

Parallelization. We ran experiments to evaluate the effect of parallelization on the verification of our model, the results of which are shown in Section 5.7. The experiments' result can serve as a justification for the number of computers used in the main presentation of the verification results, which is ten. It can be seen that although smaller number of workstations would suffice, the verification would get substantially slower in the case of larger property automata (e.g. the use-case scenarios). One of the use cases could not even be handled with only one workstation. On the other hand, larger number of workstations causes inadequate memory overhead in the case of small property automata, and the time decrease is not as substantial with more than ten workstations. The choice of ten computers seems a reasonable compromise then.

7 Conclusion and future work

In this paper, we give a practical application of the presented CI-LTL verification technique to a large component-based system using a parallel model checking tool DiVinE. We briefly introduce our modelling language as well as the temporal logic CI-LTL, a modification of the action based LTL. We have verified a multitude of properties of the Trading System. Twelve of them that are of particular interest within the component-oriented software engineering society, are presented here together with the results of the verification and their discussion. The presented properties include two basic properties describing the broadcasting ability of the event-channel components, three properties concerning the possibility of a local deadlock, two properties addressing the component blocking problem, and two properties dealing with the problems caused by cycles in the model. The last three properties are different from the previous. They are used for checking the correctness of the use-case scenarios. Finally, we discuss how the model checking helped us in creation of the model, and we summarize the experience obtained during verification, including discussion of some of the results, and the effect of the parallelization.

The study confirms that the CI automata modelling language suits well both to capture various types of interactions among individual components in component-

based systems, and to formally verify interaction properties. This distinguishes our modelling approach from others presented in the *CoCoME Modelling Contest* [15] and brings a new value to the area of component-based software engineering. As the very significant feature of component-based systems is the concurrent behaviour of individual components and consequently the enormous size of the state space, distributed and parallel verification techniques are a need for handling these type of systems in reasonable time. They allowed us to verify very complex properties of the Trading System when restricted to a usage profile. But still, we were not able to verify the Trading System with no usage profile added—this means any usage possible with any number of users—as our hardware capacity did not suffice.

In future, we aim at extending our verification techniques with various reduction methods to allow us to verify even larger systems. Currently, we explore the possibilities of two existing reduction techniques, the partial-order reduction and the symmetry reduction. However, their application in our framework is not straightforward, due to the nature of the temporal logic we use. We also try to find new reduction methods taking advantage of component-specific characteristics of verified systems.

References

[1] Adamek, J. and F. Plasil, *Behavior protocols capturing errors and updates*, in: *Proceedings of the ETAPS Workshop on Unanticipated Software Evolution (USE'03)* (2003), pp. 17–25.

[2] Allen, R. J., "A Formal Approach to Software Architecture," Ph.D. thesis, Carnegie Mellon University, School of Computer Science, USA (1997).

[3] Barnat, J., L. Brim and I. Černá, *Distributed Analysis of Large Systems*, in: *Proc. of the 4th International Symposium on Formal Methods for Components and Objects (FMCO'05)*, LNCS **2006** (2006), pp. 259–279.

[4] Barnat, J., L. Brim, I. Černá, P. Moravec, P. Ročkai and P. Šimecek, *Divine – a tool for distributed verification*, in: *Proceedings of the Computer Aided Verification conference (CAV'06)* (2006), pp. 278–281.

[5] Becker, S., H. Koziolek and R. Reussner, *Modelbased performance prediction with the palladio component model*, in: *Proceedings of the International Workshop on Software and Performance (WOSP'07)* (2007), pp. 54–65.

[6] Brim, L., I. Černá, P. Vařeková and B. Zimmerova, *Component-Interaction automata as a verification-oriented component-based system specification*, in: *Proceedings of the ESEC/FSE Workshop on Specification and Verification of Component-Based Systems (SAVCBS'05)* (2005), pp. 31–38, published also in ACM SIGSOFT Software Engineering Notes, Volume 31, Issue 2 (March 2006).

[7] Bruneton, E., T. Coupaye, M. Leclercq, V. Quéma and J.-B. Stefani, *The fractal component model and its support in java*, Software: Practice and Experience **36** (2006), pp. 1257–1284.

[8] Černá, I., P. Vařeková and B. Zimmerova, *Component-interaction automata modelling language*, Technical Report FIMU-RS-2006-08, Masaryk University, Faculty of Informatics, Brno, Czech Republic (2006).

[9] de Alfaro, L. and T. A. Henzinger, *Interface-based design*, in: *Proceedings of the 2004 Marktoberdorf Summer School* (2005), pp. 1 – 25.

[10] *DiVinE project web page*. URL http://anna.fi.muni.cz/divine/

[11] Garlan, D., R. T. Monroe and D. Wile, "Foundations of Component-Based Systems," Cambridge University Press, USA, 2000 ISBN 0-521-77164-1.

[12] Magee, J., J. Kramer and D. Giannakopoulou, *Behaviour analysis of software architectures*, in: *Proceedings of the 1st Working IFIP Conference on Software Architecture (WICSA'99)* (1999), pp. 35–50.

[13] Plasil, F. and S. Visnovsky, *Behavior protocols for software components*, IEEE Transactions on Software Engineering **28** (2002), pp. 1056–1076.

[14] Pnueli, A., *The temporal logic of programs*, in: *Proceedings of the 18th IEEE Symposium on the Foundations of Computer Science* (1977), pp. 46–57.

[15] Rausch, A., R. Reussner, R. Mirandola and F. Plasil, editors, "The Common Component Modeling Example: Comparing Software Component Models," To appear in LNCS, 2007.
URL http://www.cocome.org

[16] The CoIn Team, *The complete CoIn model of the Trading System* (2007).
URL http://anna.fi.muni.cz/coin/cocome/

[17] Vardi, M. Y., *An automata-theoretic approach to linear temporal logic*, in: *Logics for Concurrency: Structure versus Automata*, LNCS **1043** (1996), pp. 238 – 266.

[18] Zimmerova, B., P. Vařeková, N. Beneš, I. Černá, L. Brim and J. Sochor, "The Common Component Modeling Example: Comparing Software Component Models, chapter Component-Interaction Automata Approach (CoIn)," To appear in LNCS, 2007 .