

# From HELENA Ensemble Specifications to PROMELA Verification Models

Annabelle Klarl

Ludwig-Maximilians-Universität München, Germany\*

**Abstract.** With HELENA, we introduced a modeling approach for distributed systems where components dynamically collaborate in ensembles. Conceptually, components participate in a goal-oriented collaboration by adopting certain roles in the ensemble. To verify the goal-directed behavior of ensembles, we propose to systematically translate HELENA specifications to PROMELA and verify them with the model-checker Spin. In this paper, we report on tool support for an automated transition from HELENA to PROMELA. Relying on the XTEXT workbench of Eclipse, we provide a code generator from the domain-specific-language HELENA-TEXT to PROMELA. The generated PROMELA model simulates the two layers, components and their adopted roles from HELENA, and allows dynamic role creation as well as asynchronous communication of roles.

## 1 Introduction

Ensemble-based systems are distributed systems of components which dynamically collaborate in groups. In HELENA [5], components are thought of as a basic layer providing computing power or storage resources. Collaborations are modeled by *ensembles*, where components adopt (possibly concurrently) different *roles* to actively participate in ensembles. The concept of roles allows to focus on the particular tasks which components fulfill in collaborations and to structure implementation by realizing roles as threads executed on top of components [9].

Ensembles always collaborate for some global goal. Such goals are often temporal properties and are therefore specified in linear temporal logic (LTL) [11]. To allow verification of HELENA models for goals, we already proposed in [6] to translate HELENA to PROMELA and check satisfaction of goals with the model-checker Spin [7]. We proved the correctness of the translation for a simplified variant of HELENA which restricts ensemble specifications to their core concepts.

In this paper, we report on the extension of the translation to full HELENA and its automation based on the XTEXT workbench of Eclipse. With the extended translation, we are able to simulate the two layers of HELENA, components and their adopted roles, in PROMELA. Due to the automation of the translation, we augment HELENA ensemble specifications with immediate verification support in Spin. To this end, an Eclipse plug-in is implemented which produces an executable PROMELA specification from a HELENA ensemble specification written in the domain-specific language HELENA-TEXT [8].

---

\* This work has been partially sponsored by the EU project ASCENS, 257414.

## 2 HELENA in a Nutshell

We introduce the concepts of the HELENA approach at a peer-2-peer network supporting the distributed storage of files which can be retrieved upon request.

**Components:** The foundation of HELENA ensembles [5] are *components* characterized by their type, e.g., component type `Peer` in Fig. 1. Such a type manages associations to other components, e.g., the association `neighbor` in our example. It stores basic information, that is useful in all roles the component can adopt, in attributes, e.g., the attribute `hasFile`. Lastly, it provides operations which can be invoked by its roles, e.g., the operation `printFile` (not shown).

**Roles:** A *role type*  $rt$  is a tuple  $(rtnm, rtcomptypes, rtattrs, rtmsgs)$ :  $rtnm$  is the name of the role type; the set  $rtcomptypes$  determines the component types which can adopt the role; the set  $rtattrs$  allows to store data that is only relevant for performing the role; the set  $rtmsgs$  determines the outgoing and incoming messages supported by the role. In our example, we have three role types which can all be adopted by components of the type `Peer` (cf. Fig. 1): The peer adopting the role `Requester` wants to download the file, peers adopting the role `Router` forward the request through the network, and the peer adopting the role `Provider` provides the file. Only the role type `Requester` has an attribute. Outgoing and incoming messages are annotated as arrows for all role types.

**Ensemble Structures:** To define the structural characteristics of a collaboration, an *ensemble structure* specifies the role types whose instances form the ensemble, determines how many instances of each role type may contribute by a multiplicity, and defines the capacity of the input queue of each role type. We assume that between two role types the messages which are output on one side and input on the other side can be exchanged. For our example, instances of the three aforementioned role types collaborate (cf. Fig. 1). Thereby, an ensemble has to employ exactly one requesting peer, arbitrarily many routers, and possibly one router as determined by the multiplicities associated to each role type.

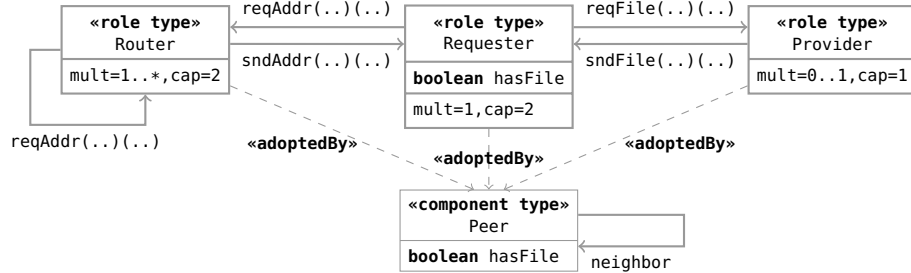


Fig. 1: Ensemble structure for the p2p example in graphical notation

**Ensemble Specifications:** The *behavior* of a role is specified by a process expression built from the null process `nil`, action prefix  $a.P$ , guarded choice  $\mathbf{if}(guard_1) \{P_1\} \mathbf{or}(guard_2) \{P_2\}$  (branch is nondeterministically selected if several branches are executable), and process invocation [6]. Guards are predicates over component or role attributes. There are actions for creating (**create**) and

retrieving (**get**) role instances, sending (!) or receiving (?) messages, and invoking operations of the owning component. These actions must fit to the declared ensemble structure, e.g., messages can be only sent by roles which declare them. Additionally, state labels are used to mark a certain progress of execution in the role behavior. Fig. 2 shows the behavior specification of a **Router**. Initially, a router can receive a request for an address. Depending on whether its owner has the file, it either creates a provider role instance and sends it back to the requester in  $P_{provide}$  or forwards the request to another router in  $P_{fwd}$  if possible.

```

roleBehavior Router = ?reqAddr(Requester rq)() .
    if (owner.hasFile) then {Pprovide}
    or (!owner.hasFile) then {Pfwd}
    Pprovide = p←create(Provider, owner) . rq!sndAddr(p)() . nil
    Pfwd = if (plays(Router, owner.neighbor)) then {nil}
    or (!plays(Router, owner.neighbor)) then {Pcreate}
    Pcreate = r←create(Router, owner.neighbor) . r!reqAddr(rq)() . Router

```

Fig. 2: Role behavior of a **Router** for the p2p example

A complete collaboration is given by an *ensemble specification* consisting of an ensemble structure  $\Sigma$  and a set of role behaviors, one for each role type in  $\Sigma$ . The complete specification of the example can be found in [10].

**Semantics:** Ensemble specifications are semantically interpreted by labeled transition systems, i.e., ensemble automata [5,6]. Ensemble states capture the currently existing role instances with their data and control states. Transitions between ensemble states are triggered by role instance creation or retrieval, communication actions, and operation calls. The communication style (synchronous or asynchronous) is determined by the size of the input queues of the role types.

**Goal Specifications:** Goals are expressed by LTL formulae over particular HELENA propositions: A *state label proposition* is of the form  $rt[n]@label$ . It is satisfied if there exists a role instance  $n$  of type  $rt$  whose next performed action is the state label  $label$ . An *attribute proposition* must be boolean and is built from arithmetic and relational operators, data constants, and propositions of the form  $rt[n]:attr$  (or  $ct[n]:attr$ ). An attribute proposition  $rt[n]:attr$  is satisfied if there exists a role instance  $n$  of type  $rt$  such that the value of its attribute  $attr$  evaluates to **true** (and analogously for component attributes). LTL formulae and their satisfaction are inductively defined from HELENA propositions, propositional operators  $\neg$  and  $\wedge$  and LTL operators **X**,  $\square$ ,  $\diamond$ , **U** and **W** as usual.

For the p2p example, we want to express that the requester will always receive the requested file if the file is available in the network. We assume a network of three peers and formulate the following achieve goal in LTL which refers to the values of the attribute `hasFile` of component type **Peer** and role type **Requester**:

$$(Peer[1]:hasFile \vee Peer[2]:hasFile \vee Peer[3]:hasFile) \Rightarrow \diamond Requester[1]:hasFile$$

### 3 Translation from HELENA to PROMELA

To verify HELENA specifications for their intended goals, we rely on the model-checker Spin [7]. In [6], we discussed that the translation of a simplified variant

of HELENA to PROMELA preserves satisfaction of  $LTL_{\setminus \mathbf{X}}$ , the fragment of LTL that does not contain the *next* operator  $\mathbf{X}$ . This translation abstracts from the underlying component-based platform and considers only role types and their interactions. In role behaviors, guarded choice and arbitrary process invocations are not allowed and any notion of data is omitted. To cope with these features, we propose to represent components and roles by two kinds of processes in PROMELA. They differ in communication abilities and behavior since components are only storage and computing resources while roles are active entities.

**Communication Abilities:** (1) Components only interact with roles, but not with other components. Roles advise components to adopt other roles, request references to already adopted roles from their owning components, or invoke operations on them. Thus, each PROMELA process for a component relies on a dedicated synchronous channel **self**, only used for communication between itself and its adopted roles. The roles refer to the channel under the name **owner**. (2) Roles interact by exchanging directed messages on input queues. Thus, each PROMELA process for a role relies on a dedicated (possibly asynchronous) channel **self** in addition to the aforementioned channel **owner** to model its input queue. Since channels are global in PROMELA, but input queues are local in HELENA, special care has to be taken that this channel is only available to processes which are allowed to communicate with the corresponding role in HELENA.

**Behavior:** (1) The PROMELA process for a component implements a **do**-loop to wait for requests from its roles on the **self** channel. Depending on the request, it runs some internal computation and sends a reply. E.g., to adopt a role, it creates a new channel and spawns a new process (representing the role) to which it hands over its own **self** channel as the role's **owner** channel and the newly created channel as the role's **self** channel. Afterwards, it sends the role's **self** channel to the role requesting the adoption such that the two roles can communicate via this channel. (2) The HELENA role behaviors must be reflected by the corresponding PROMELA process. In [6], we proposed to translate action prefix to sequential composition, nondeterministic choice to the **if**-construct, and recursive behavior invocation to a **goto** to the beginning of the role behavior. Sending and receiving messages was mapped to message exchange on the **self** channel of roles and role creation to process creation with the **run**-command. To extend this to full HELENA, guarded choice is translated to the **if**-construct with the guard as first statement. Arbitrary process invocation is realized by jumping to labels marking the beginning of processes. On the level of actions, we extend message exchange by data relying on user-defined data types in PROMELA. To cope with the component level of HELENA, a new role is created by issuing an appropriate request on the owning component and spawning the new role process from there. The introduction of components also allows us to implement role retrieval and operation calls by corresponding requests from role to component.

**$LTL_{\setminus \mathbf{X}}$  Preservation:** Similarly to the simplified translation in [6], all HELENA constructs are directly translated to PROMELA while introducing some additional silent steps like **gotos**. These do not hamper stutter trace equivalence and thus satisfaction of  $LTL_{\setminus \mathbf{X}}$  is preserved, though not formally shown here.

## 4 Automation of the Translation

To automate the translation, a code generator, taking a HELENATEXT [8] ensemble specification as input, was implemented on top of the XTEXT workbench of Eclipse relying on XTEND as a template language.

**Component Types:** For each component type, the excerpt of the XTEND template in Fig. 3 generates a new process type in PROMELA. Most importantly, this process type implements a **do**-loop (line 4-10) where it can repeatedly receive requests from its adopted roles via its **self** channel. Depending on the type of the received request, i.e., `req.optype`, it either executes an operation (line 7), adopts a new role (line 9), or retrieves an already existing one (line 10).

```
1 def static compileProctype(ComponentType ct, Iterable<RoleType> roleTypes) {
2   ''' proctype <ct.name>(chan self; ...) {
3     <<FOR rt:roleTypes>> chan <rt.name> = [<rt.capacity>] of { Msg }; ...
4     do
5       ::self?req ->
6         if
7           <<FOR o:ct.ops>> ::req.optype==<o.name> -> // execute operation ...
8           <<FOR rt:roleTypes>>
9             ::req.optype==<rt.create>-> ...run <rt.name>(self,<rt.name>);answer!<rt.name>
10            ::req.optype==<rt.get> -> ...answer!<rt.name> ...
```

Fig. 3: Excerpt of the XTEND template for the translation of component types

**Role Types:** For each role type, the XTEND template in Fig. 4 generates a new process type in PROMELA. Two parameters for the **owner** and **self** channels are declared (line 2) and the role behavior is translated (line 3), e.g., action prefix is represented by sequential composition (line 4-6) and guarded choice by an **if**-construct (line 7-14). Furthermore, the generation of the reception of messages and create actions is shown in the right part of Fig. 4 since they represent two different types of communication: An incoming message is represented by a user-defined data type `Msg` (line 2), to cope with data parameters, and is received on the **self** channel (line 3). The role checks whether the received message was actually expected (line 4) and unpacks its parameters (line 5-7). For a create action, the component `crt.comp` is asked to adopt a role of type `crt.roleInst.type` (line 12). The component is responsible for creating the role (cf. Fig. 3) and sends back the **self** channel of the newly created role (line 13). The implementation of the generator and the HELENATEXT specification of the p2p example as well as its generated PROMELA translation can be found in [10].

## 5 Conclusion

We presented how to verify HELENA specifications for goals specified by LTL formulae with the model-checker Spin. We defined a translation of HELENA specifications and its two-layered architecture into PROMELA which was implemented on top of XTEXT. In first experiments with larger case studies, the application of Spin scales well with the size of the HELENA model since the state space only grows by a constant factor compared to HELENA. For future work, we especially want to add support for relating the Spin output back to HELENA.

```

1 def genRoleBehavior(RoleBehavior rb) {
2   ''' proctype «rb.name»(chan owner,self){
3     «rb.genProcTerm» ...
4   def genProcTerm(ActionPrefix term) {
5     ''' «term.action.genAction»;
6     «term.procTerm.genProcTerm» ...
7   def genProcTerm(GuardedChoice term) {
8     ''' if
9     ::(«term.ifGuard.genGuard») ->
10    «term.ifProcTerm.genProcTerm»
11    «FOR i : 0 ..< term.orGuards.size»
12    ::(«term.orGuards.get(i).genGuard») ->
13    «term.orProcTerms.get(i).genProcTerm»
14  }
15 }
16 }

```

```

1 def genAction(IncomingMessage m) {
2   ''' Msg «m.name»;
3   self?«m.name»;
4   «m.name».msgtype == «m.type»;
5   «FOR p:m.rparams» chan «p.name» = ...;
6   «FOR p:m.dparams»
7     «p.type» «p.name» = ...;
8   ...
9   def genAction(CreateAction crt) {
10    ''' chan «crt.roleInst.name»;
11    chan answer = [0] of { chan };
12    «crt.comp!»«crt.roleInst.type»,answer;
13    answer?«crt.roleInst.name»;
14  }
15 }

```

Fig. 4: Excerpt of the XTEND template for the translation of role types

Our approach is in-line with the goal-oriented requirements approach KAOS [11]. However, KAOS specifications are translated to the process algebra FSP which cannot represent directed communication and dynamic process creation. Furthermore, techniques for verifying ensemble-based systems have been proposed. In [4], ensembles are described by simplified SCEL programs and translated to PROMELA, but the translation is neither proved correct nor automated and cannot cope with dynamic creation of components. DFINDER [2] implements efficient strategies exploiting compositional verification of invariants to prove safety properties for BIP ensembles, but again does not deal with dynamic creation of components. DEECo ensembles [1] are implemented with the Java framework jDEECo and verified with Java Pathfinder [2]. Thus, they do not need any translation. However, since DEECo relies on knowledge exchange rather than message passing, they do not verify any communication behaviors.

## References

1. Bures, T., Gerostathopoulos, I., Hnetyuka, P., Keznikl, J., Kit, M.: DEECO: An Ensemble-based Component System. In: CBSE 2013. pp. 81–90. ACM (2013)
2. Combaz, J., Bensalem, S., Kofron, J.: Correctness of Service Components and Service Component Ensembles. In: Software Engineering for Collective Autonomic Systems, LNCS, vol. 8998. Springer (2015)
3. De Nicola, R., Latella, D., Lafuente, A.L., Loreti, M., Margheri, A., Massink, M.: The SCEL Language: Design, Implementation, Verification. In: Software Engineering for Collective Autonomic Systems. LNCS, vol. 8998. Springer (2015)
4. De Nicola, R., Lluch-Lafuente, A., Loreti, M., Morichetta, A., Pugliese, R., Senni, V., Tiezzi, F.: Programming and Verifying Component Ensembles. In: From Programs to Systems. LNCS, vol. 8415, pp. 69–83. Springer (2014)
5. Hennicker, R., Klarl, A.: Foundations for Ensemble Modeling - The Helena Approach. In: SAS 2014. LNCS, vol. 8373, pp. 359–381. Springer (2014)
6. Hennicker, R., Klarl, A., Wirsing, M.: Model-Checking Helena Specifications with Spin. In: LRC 2015. LNCS, Springer (to appear 2015), <http://goo.gl/aldya2>
7. Holzmann, G.: The Spin Model Checker. Addison-Wesley (2003)
8. Klarl, A., Cichella, L., Hennicker, R.: From Helena Ensemble Specifications to Executable Code. In: FACS 2014. LNCS, vol. 8997, pp. 183–190. Springer (2014)
9. Klarl, A., Hennicker, R.: Design and Implementation of Dynamically Evolving Ensembles with the Helena Framework. In: ASWEC 2014. pp. 15–24. IEEE (2014)
10. Klarl, A., Hennicker, R.: The Helena Framework (2015), <http://goo.gl/aldya2>
11. van Lamsweerde, A.: Requirements Engineering: From System Goals to UML Models to Software Specifications. Wiley (2009)