

The ADORA Approach to Object-Oriented Modeling of Software

Martin Glinz¹, Stefan Berner², Stefan Joos³, Johannes Ryser¹,
Nancy Schett¹, and Yong Xia¹

¹ Institut für Informatik, Universität Zürich, Winterthurerstrasse 190,
CH-8057 Zurich, Switzerland

{glinz, ryser, schett, xia}@ifi.unizh.ch

² FJA, Zollikerstrasse 183, CH-8008 Zurich, Switzerland

stefan.berner@fja.com

³ Robert Bosch GmbH, Postfach 30 02 20, D-70469 Stuttgart, Germany

stefan.joos@de.bosch.com

Abstract. In this paper, we present the ADORA approach to object-oriented modeling of software (ADORA stands for Analysis and Description of Requirements and Architecture). The main features of ADORA that distinguish it from other approaches like UML are the *use of abstract objects* (instead of classes) as the basis of the model, a *systematic hierarchical decomposition* of the modeled system and the *integration* of all aspects of the system *in one coherent model*. The paper introduces the concepts of ADORA and the rationale behind them, gives an overview of the language, and reports the results of a validation experiment for the ADORA language.

1 Introduction

When we started our work on object-oriented specification some years ago, we were motivated by the severe weaknesses of the then existing methods, e.g. [3][5][15]. In the meantime, the advent of UML [16] (and to a minor extent, OML [6]) has radically changed the landscape of object-oriented specification languages. However, also with UML and OML, several major problems remain.

There is still no true integration of the aspects of data, functionality, behavior and user interaction. Neither do we have a systematic hierarchical decomposition of models (for example, UML packages are a simple container construct with nearly no semantics). Models of system context and of user-oriented external behavior are weak and badly integrated with the class/object model [9].

So there is still enough motivation not to join simply the UML mainstream and to pursue alternatives instead. We are developing an object-oriented modeling method for software that we call ADORA (Analysis and Description of Requirements and Architecture) [1][13]. ADORA is intended to be used primarily for requirements specification and also for logical-level architectural design. Currently, ADORA has no language elements for expressing physical design models (distribution, deployment) and imple-

mentation models. However, as ADORA models are object-oriented, we can implement a smooth transition from an ADORA architecture model to detailed design and code written in an object-oriented programming language.

In this paper, we present the ADORA language. We discuss the general concepts and give an overview of the language. The main contributions of the ADORA language are

- a concept for systematic hierarchical decomposition of models which is particularly useful when modeling distributed systems,
- the integration of different aspects into one coherent model,
- the ability to visualize a model in its context,
- language elements for tailoring the formality of ADORA models.

Throughout this paper, we will use a distributed heating control system as an example. The goal of this system is to provide a comfortable control for the heating system of a building with several rooms. An operator can control the complete system, setting default temperatures for the rooms. Additionally, for every room individual temperature control can be enabled by the operator. Users then can set the desired temperature using a control panel in the room. The system shall be distributed into one master module serving the operator and many room modules.

However, ADORA is not only applicable for the specification of industrial control systems. For the validation of the usefulness of the language, we have modeled a distributed information system with ADORA (see section 4).

The rest of the paper is organized as follows. In section 2 we discuss the basic concepts of ADORA and their rationale. In section 3 we give an overview of the language. In section 4 we present the results of a first validation of the ADORA language. Finally, we compare the concepts of ADORA with those of UML and conclude with a discussion of results, state of work and future directions.

2 Key Concepts and Rationale of the ADORA Approach

In this section, we briefly describe the five principles that ADORA is based on and give our rationale for choosing them.

2.1 Abstract Objects instead of Classes

When we started the ADORA project, all existing object-oriented modeling methods used class diagrams as their model cornerstone. However, class models are inappropriate when more than one object of the same class and/or collaboration between objects have to be modeled [9][12]. Both situations frequently occur in practice. For an example, see the buttons in Fig. 1. Moreover, class models are difficult to decompose. As soon as different objects of a class belong to different parts of a system (which often is the case), hierarchical decomposition does no longer work for class models [12]. Wirfs-Brock [19] tries to overcome the problems of class modeling by using classes in

different *roles*. However, decomposition remains a problem: what does it mean to decompose a role?

We therefore decided to use abstract, prototypical objects as the core of an ADORA model (Fig. 1). An equivalent to classes (which we call types) is only used to model common characteristics of objects: types define the properties of the objects and can be organized in subtype hierarchies. In order to make models more precise, we distinguish between *objects* (representing a single instance) and *object sets* that represent a set of instances. Modeling of collaboration and of hierarchical decomposition (see below) becomes easy and straightforward with abstract objects.

In the meantime, others have also discovered the benefits of modeling with abstract objects. UML, for example, uses abstract objects for modeling collaboration in collaboration diagrams and in sequence diagrams. However, without a notion of abstraction and decomposition, only local views can be modeled. Moreover, class diagrams still form the core of a UML specification.

2.2 Hierarchical Decomposition

Every large specification must be decomposed in some way in order to make it manageable and comprehensible. A good decomposition (one that follows the basic software engineering principles of information hiding and separation of concerns) decomposes a system recursively into parts such that

- every part is logically coherent, shares information with other parts only through narrow interfaces and can be understood in detail without detailed knowledge of other parts,
- every composite gives an abstract overview of its parts and their interrelationships.

The current object-oriented modeling methods typically approach the decomposition problem in two ways: (a) by modeling systems as collections of models where each model represents a different aspect or gives a partial view of the system, and (b) by providing a container construct in the language that allows the modeler to partition a model into chunks of related information (e.g. packages in UML). However, both ways do not satisfy the criteria of a good decomposition. Aspect and view decompositions are coherent only as far as the particular aspect or view is concerned. The infor-

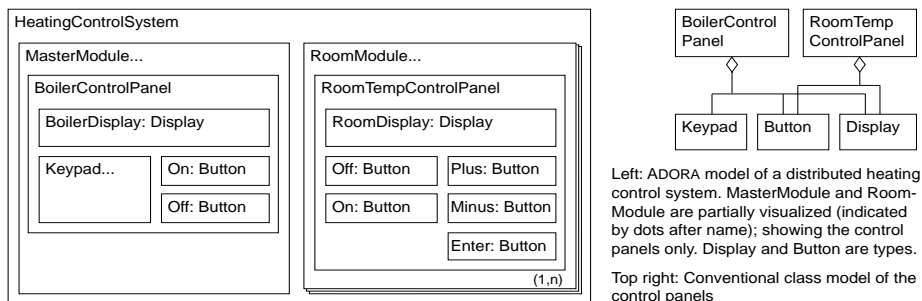


Fig. 1. An ADORA object model (left) vs. a conventional class model (top right)

mation required for comprehending some part of a system in detail is *not* coherently provided. Container constructs such as UML packages have semantics that are too weak for serving as composites in the sense that the composite is an abstract overview of its parts and their interrelationships. This is particularly true for multi-level decompositions. Only the ROOM method [18] can decompose a system in a systematic way. However, as ROOM is also based on classes, the components are not classes, but class references. This asymmetry makes it impossible to define multi-level decompositions in a straightforward, easily understandable way.

In ADORA, the decomposition mechanism was deliberately chosen such that good decompositions in the sense of the definition given above become possible. We recursively decompose objects into objects (or elements that may be part of an object, like states). So we have the full power of object modeling on all levels of the hierarchy and only vary the degree of abstractness: objects on lower levels of the decomposition model small parts of a system in detail, whereas objects on higher levels model large parts or the whole system on an abstract level.

2.3 Integrated Model

With existing modeling languages, one creates models that consist of a set of more or less loosely coupled diagrams of different types. UML is the most prominent example of this style. This seems to be a good way to achieve separation of concerns. However, while making separation easy, loosely coupled collections of models make the equally important issues of integration and abstraction of concerns quite difficult.

In contrast to the approach of UML and others, an ADORA model integrates all modeling aspects (structure, data, behavior, user interaction...) in one coherent model. This allows us to develop a strong notion of consistency and provides the necessary basis for developing powerful consistency checking mechanisms in tools. Moreover, an integrated model makes model construction more systematic, reduces redundancy and simplifies completeness checking.

Using an integrated model does of course not mean that everything is drawn in one single diagram. Doing so would drown the user in a flood of information. We achieve separation of concerns in two ways: (1) We *decompose* the model *hierarchically*, thus allowing the user to select the focus and the level of abstraction. (2) We use a *view concept* that is based on *aspects*, not on various diagram types. The *base view* consists of the objects and their hierarchical structure only. The base view *is combined with* one or more *aspect views*, depending on what the user wishes to see. These two concepts – hierarchy and combination of views – constitute the *essence* of organizing an ADORA model.

So the complete model is basically an abstract one – it is almost never drawn in a diagram. The concrete diagrams typically illustrate certain aspects of certain parts of a model in their hierarchical context. However, as every concrete diagram is a view of an integrated model of the complete system, we can build strong consistency and completeness rules into the language and build powerful tools for checking and maintaining them. Readability of diagrams is achieved by selecting the right level of abstraction, by restricting the number of aspects being viewed together, and by split-

ting complex diagrams into an abstract overview diagram and some part diagrams. For example, if Fig. 2 is perceived to be too complex, it can be split into an overview diagram (Fig. 8) and two part diagrams, one for `MasterModule` and one for `RoomModule`.

2.4 Adaptable Degree of Formality

An industrial-scale modeling language should allow its users to adapt the degree of formalism in a specification to the difficulty and the risk of the problem in hand. Therefore, they need a language with a broad spectrum of formality in its constructs, ranging from natural language to completely formal elements.

In ADORA, we satisfy this requirement by giving the modeler a choice between informal, textual specifications and formal specifications (or a mixture of both). For example, an object may be specified with an informal text only. Alternatively, it can be formally decomposed into components. These in turn can be specified formally or informally. As another example, state transitions can be specified in a formal notation or informally with text or with a combination of both.

The syntax of the ADORA language provides a consistent framework for the use of constructs with different degrees of formality.

2.5 Contextual Visualization

Current modeling languages either lack capabilities for system decomposition or they visualize decompositions in an explosive zoom style: the composites and their parts are drawn as separate diagrams. Thus, a diagram gives no information about the context that the presented model elements are embedded in. In ADORA, we use a fisheye view concept for visualizing a component in its surrounding hierarchical context. This simplifies browsing through a set of diagrams and improves comprehensibility [1].

3 An Overview of the ADORA Language

An ADORA model consists of a basic hierarchical object structure (the base view, as we call it) and a set of aspect views that are combined with the base view. In this section we describe these views and their interaction.

3.1 Basic Hierarchical Object Structure

The object hierarchy forms the basic structure of an ADORA model.

Objects and object sets. As already mentioned above, we distinguish between objects and object sets. An *ADORA object* is an abstract representation of a single instance in the system being modeled. For example, in our heating control system, there is a single boiler control panel, so we model this entity as an object. Abstract means that the object is a placeholder for a concrete object instance. While every object instance must

have an object identifier and concrete values for its attributes, an ADORA object has neither of these. An *ADORA object set* is an abstract representation of a set of object instances. The number of instances allowed can be constrained with a cardinality. For example, in an order processing system we would model suppliers, parts, orders, etc. as object sets. In the heating control system, we have a control panel in every room and we control at least one room. Thus we model this panel as an object set with cardinality (1,n); see Fig. 2.

Structure of an ADORA object. An object or object set has a twofold inner structure: it consists of a set of properties and (optionally) a set of parts.

The *properties* are attributes (both public and private ones), directed relationships to other objects/object sets, operations and so called standardized properties. The latter are user-definable structures for stating goals, constraints, configuration information, notes, etc.; see below.

The *parts* can be objects, object sets, states and scenarios. Every part again can consist of parts: objects and object sets can be decomposed recursively as defined above, states can be refined into statecharts, scenarios into scenariocharts (as we call them, see below). Thus, we get a hierarchical whole-part structure that allows modeling a hierarchical decomposition of a system. The decomposition is strict: an element neither can contain itself nor can it be a part of more than one composite. We stick to a strict decomposition due to its inherent simplicity and elegance. Commonalities between objects in different positions of a decomposition hierarchy can be modeled by assigning them the same type (see the paragraph on types below and Fig. 1).

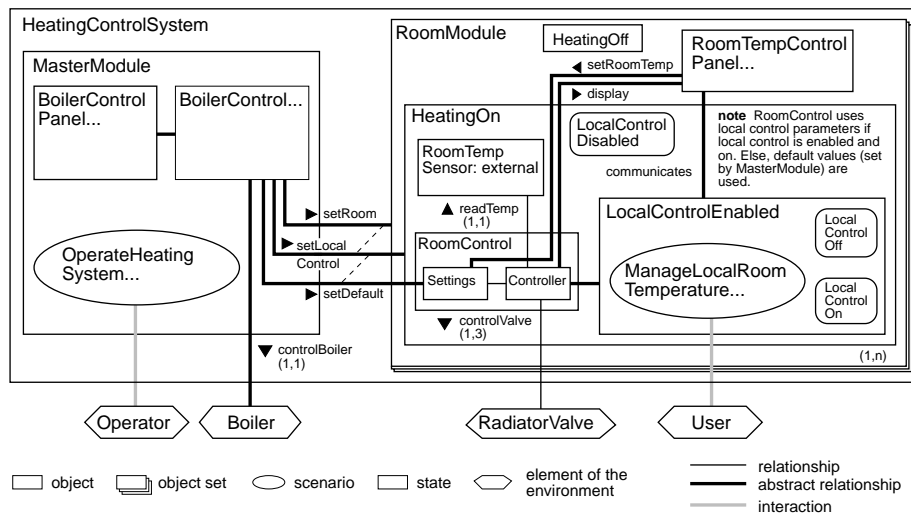


Fig. 2. An ADORA view of the heating control system: base view combined with structural view and context view

Graphic representation. In order to exploit the power of hierarchical decomposition, we allow the modelers to represent an ADORA model on any level of abstraction, from a very high-level view of the complete system down to very detailed views of its parts.

We achieve this property by representing ADORA objects, object sets, scenarios and states by nested boxes (see Fig. 1 and 2). The modeler can freely choose between drawing few diagrams with deep nesting and more diagrams with little nesting. In order to distinguish expanded and non-expanded elements in a diagram, we append three dots to the name of every element having parts that are not or only partially drawn on that diagram.

Types. Frequently, different objects have the same inner structure, but are embedded in different parts of a system. In the heating system for example, the boiler control panel and the room control panels both might have a display with the same properties. In these situations, it would be cumbersome to define the properties individually for every object. Instead, ADORA offers a type construct. An ADORA type defines

- the attributes and operations of all objects/object sets of this type
- a structural interface, that means, information required from or provided to the environment of any object/object set of this type. This facility can be used to express *contracts*.

A type defines neither the relationships to other objects/object sets nor the embedding of the objects of that type. Types can be organized in a subtype hierarchy.

An object can have a type name appended to its name (for example, RoomDisplay: Display in Fig. 1). In this case, the object is of that type and the type is separately defined in textual form. Otherwise, there is no other object of the same type in the model and the type information is included in the definition of the object.

```
propertydef goal numbered Hyperstring constraints unique;
propertydef created Date;
propertydef note Hyperstring;
object HeatingControlSystem...
goal 1 "Provide a comfortable control for the heating of a building with several rooms."
created 2000-11-04
note "Constraints have yet to be discussed and added."
end HeatingControlSystem.
```

Fig. 3. Definition and use of standardized properties

Standardized properties. In order to adapt ADORA in a flexible, yet controlled way to the needs of different projects, application domains or persons, we provide so called *standardized properties*. An ADORA standardized property is a typed construct consisting of a header and a body. Fig. 3 shows the type definitions for the properties goal, created and note and the application of these properties in the specification of the object HeatingControlSystem. As name and structure of the properties are user-definable, we get the required flexibility. On the other hand, typing ensures that a tool nevertheless can check the properties and support searching, hyperlinking and cross-referencing.

3.2 The Structural View

The structural view combines the base view with directed relationships between objects. Whenever an object A references an information in another object B (and B is

not a part of A or vice-versa) then there must be a relationship from A to B. Referencing an information means that A

- accesses a public attribute of B,
- invokes an operation of B,
- sends an event to B or receives one from B.

Every relationship has a name and a cardinality (in the direction of the relationship). Bidirectional relationships are modeled by two names and cardinalities. Relationships are graphically represented by lines between the linked objects/object sets. An arrow preceding the name indicates the direction of the relationship (Fig. 2).


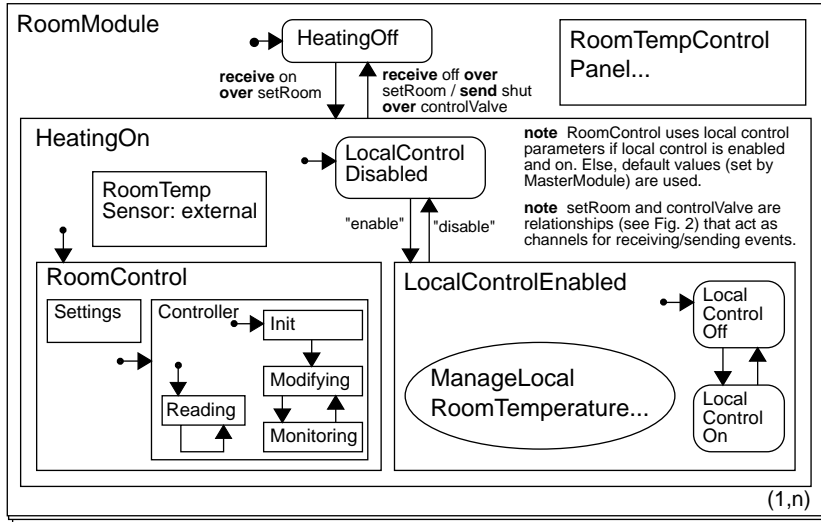


Fig. 4. Four static views of the same model on different levels of abstraction

Static relationships must reflect the hierarchical structure of the model. Let objects A and B be linked by a relationship. If A is contained in another object X and B in an object Y, then the relationship $A \rightarrow B$ implies abstract relationships $X \rightarrow Y$, $A \rightarrow Y$ and $X \rightarrow B$. Whenever we draw a diagram that hides A, B or both, the next higher abstract relationship must be drawn. Abstract relationships are drawn as thick lines. They can, but need not be named. In case of partially expanded objects, we sometimes have to draw both a concrete and a corresponding abstract relationship. In this case, we indicate the correspondence by a dashed hairline (Fig. 4). In the view shown in Fig. 2, we have some examples. All relationships from BoilerControl to other objects are abstract ones because their origins within BoilerControl are hidden in this view. The relationships readTemp from Controller to RoomTempSensor and controlValve from Controller to RadiatorValve are elementary relationships. Hence they are drawn with thin lines. If we had chosen a view that hides the contents of RoomControl, we had drawn two abstract relationships from RoomControl to RoomTempSensor and to RadiatorValve, respectively.

3.3 The Behavioral View

Combining objects and states. For modeling behavior, ADORA combines the object hierarchy with a statechart-like state machine hierarchy [7][8][12]. Every object represents an abstract state that can be refined by the objects and/or the states that an object contains. This is completely analogous to the refinement hierarchy in statecharts [10] and can be given analogous semantics for state transitions. We distinguish pure states (represented graphically by rounded rectangles) and objects with state (see Fig. 5). Pure states are either elementary or are refined by a pure statechart. Objects with state additionally have properties and/or parts other than states.



State Transition Tables for Controller

Init, Monitoring → Modifying

Modifying → Monitoring

180 s IN Modifying	Y
--------------------	---

Reading → Reading

10 s IN Reading	Y
self.ReadSensorValue	✓

Init, Monitoring → Modifying

IN LocalControlEnabled	Y	•	Y	•
IN LocalControlEnabled.LocalControlOn	Y	N	Y	N
ActualTemp > Settings.CurrentTemp(now)	N	•	Y	•
ActualTemp < Settings.CurrentTemp(now)	Y	•	N	•
ActualTemp > Settings.DefaultTemp(now)	•	N	•	Y
ActualTemp < Settings.DefaultTemp(now)	•	Y	•	N
send open over controlValve	✓	✓		
send close over controlValve			✓	✓

Fig. 5. A partial ADORA model of the heating control system; base view and behavior view

We do not explicitly separate parallel states/state machines as it is done in state-charts. Instead, objects and states that are part of the same object and have no state transitions between each other are considered to be parallel states. Objects that neither are the destination of a state transition nor are designated as initial abstract states are considered to have no explicitly modeled state.

By embedding the behavior model into the object decomposition hierarchy, we can easily model behavior on all levels of abstraction. On a high level, objects and states may represent abstract concepts like operational modes (off, startup, operating...). On the level of elementary objects, states and transitions model object life cycles.

State transitions. Triggering events and triggered actions or events can either be written in the traditional way as an adornment of the state transition arrows in the diagrams, or they can be expressed with transition tables [14]. For large systems with complex transition conditions the latter notation is more or less mandatory in order to keep the model readable. Depending on the degree of required precision, state transition expressions can be formulated textually, formally, or with a combination of both.

Fig. 5 shows the graphic representation of a behavior view with some of the variants described above. When the system is started, then for all members of the object set

RoomModule the initial state HeatingOff is entered. The transition to the object HeatingOn is specified formally. It is taken when the event on is received over the relationship set-Room (cf. Fig. 2). If this transition is taken, the state LocalControlDisabled and the object RoomControl are entered concurrently. Within RoomControl, the object Controller is entered and within Controller the parallel states Init and Reading. This is equivalent to the rules that we have for statecharts. The state transitions between LocalControlDisabled and LocalControlEnabled are specified informally with a text only. This makes sense in situations where we do not yet precisely know in which situation an event has to be triggered by which component. The transitions within Controller are specified in tabular form, because they are quite complex.

Timing and event propagation. In ADORA we use the quasi-synchronous timing and event propagation semantics defined in [8]. In contrast to usual statecharts and other than in [8] we do not broadcast events in ADORA. Instead, events have to be explicitly sent and received. Thus we avoid global propagation of local events.

3.4 The Functional View

The functional view defines the properties of an object or object set (attributes, operations...) that have not already been defined by the object's type. When there is only one object of a certain type, the complete type information is embedded in the object definition. The functional view is not combined with other views; it is always represented separately in textual form.

Joos [13] has defined a formal notation for specifying functions in ADORA, building upon existing notations. As there is nothing conceptually new with function definitions, we do not go into further detail here.

3.5 The User View

In the last few years, the importance of modeling systems from a user's viewpoint, using scenarios or use cases, was recognized (for example, see [4][8][11] and many others). In ADORA, we take the idea of hierarchically structured scenarios from [8] a step further and integrate the scenarios into the overall hierarchical structure of the system. In our terminology, a scenario is an ordered set of interactions between partners, usually between a system and a set of actors external to the system. It may comprise a concrete sequence of interaction steps (instance scenario) or a set of possible interaction steps (type scenario). Hence, a use case is a type scenario in our terminology.

We view scenarios and objects to be complementary in a specification. The scenarios specify the stimuli that actors send to the system and the system's reactions to these stimuli. However, when these reactions depend on the history of previous stimuli and reactions, that means on stored information, a precise specification of reactions with scenarios alone becomes infeasible. The objects specify the structure, functions and behavior that are needed to specify the reactions in the scenarios properly.

In the base view of an ADORA model, scenarios are represented with ovals. In the user view, we combine the base view with grey lines that link the scenario with all

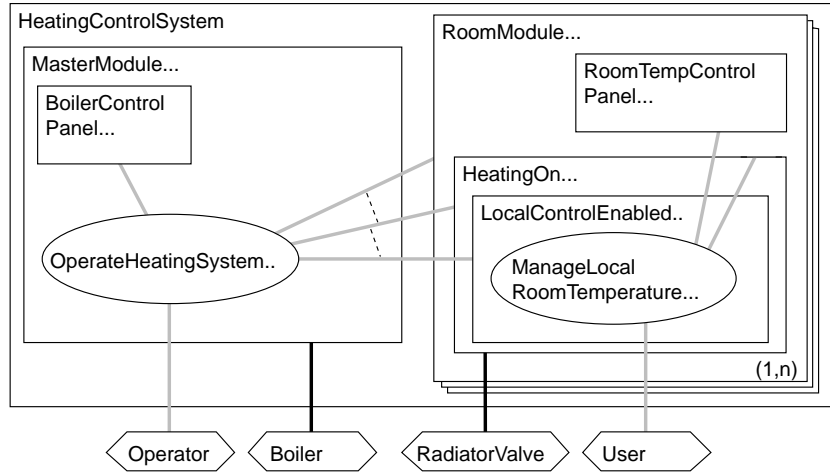


Fig. 6. A user view of the heating control system

objects that it interacts with (Fig. 6). For example, the scenario `ManageLocalRoomTemperature` specifying the interaction between the actor `User` and the system is localized within the object `LocalControlEnabled`. Internally, the scenario interacts with `RoomTempControlPanel` and with an object in `HeatingOn` which is hidden in this view.

An individual scenario can be specified textually or with a statechart. In both cases, ADORA requires scenarios to have one starting and one exit point. Thus, complex scenarios can be easily built from elementary ones, using the well-known sequence, alternative, iteration and concurrency constructors. In [8] we have demonstrated statechart-based integration of scenarios using these constructors. However, when integrating many scenarios, the resulting statechart becomes difficult to read. We therefore use Jackson-style diagrams (with a straightforward extension to include concurrency) as an additional means for visualizing scenario composition. We call these diagrams scenariocharts (Fig. 7).

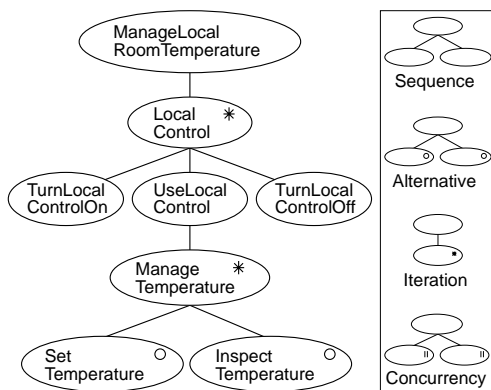


Fig. 7. A scenariochart modeling the structure of the `ManageLocalRoomTemperature` scenario

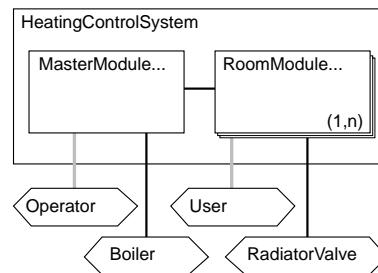


Fig. 8. A context diagram of the heating control system

Thus, we have a hierarchical decomposition in the user view, too. The object hierarchy of the base view allocates high-level scenarios (like `ManageLocalRoomTemperature` in our heating system) to that part of the system where they take effect. The scenario hierarchy decomposes high-level scenarios into more elementary ones. As a large system has a large number of scenarios (we mean type scenarios/use cases here, not instance scenarios), this facility is very important for grouping and structuring the scenarios.

Note that the ADORA user view is a logical view of user-system interaction only; it does not include the design of the user interface.

3.6 The Context View

The context view shows all actors and objects in the environment of the modeled system and their relationships with the system. Depending on the degree of abstraction selected for the system, we get a context diagram (Fig. 8) or the external context for a more detailed view of the system (Fig. 2).

In addition to external elements that are not a part of the system being specified, an ADORA model can also contain so called *external objects*. We use these to model preexisting components that are part of the system, but not part of the specification (because they already exist). External objects are treated as black boxes having a name only. In the notation, such objects are marked with the keyword `external` (for example, the object `RoomTempSensor` in Fig. 2). In any specification where COTS components will be part of the system or where existing components will be reused, modeling the embedding of these components into the system requires external objects.

3.7 Modeling Constraints and Qualities

Constraints and quality requirements are typically expressed with text, even in specifications that employ graphical models for functional specifications. In traditional specifications and with UML-style graphic models we have the problem of interrelating functional and non-functional specifications and of expressing the non-functional specifications on the right level of abstraction.

In ADORA, we use two ADORA-specific features to solve this problem. (1) The decomposition hierarchy in ADORA models is used to put every non-functional requirement into its right place. It is positioned in the hierarchy according to the scope of the requirement. The requirements themselves are expressed as ADORA standardized properties. Every kind of non-functional requirement can be expressed by its own property kind, for example performance constraint, accuracy constraint, quality...).

3.8 Consistency Checking and Model Verification

Having an integrated model allows us to define stringent rules for consistency between views, for example “When an object A references information in another object B in

any view and B is not a part of A or vice-versa, then there must be a relationship from A to B in the static view.” A language for the formulation of consistency constraints and a compiler that translates these constraints into Java have been developed [17]. By executing this code in the ADORA tool, the tool is enabled to check or enforce these constraints. The capabilities for formal analysis and verification of an ADORA model depend on the chosen degree of formality. In the behavior view, for example, a sufficiently formal specification of state transitions allows to apply all analyses that are available for hierarchical state machines.

4 Validation of ADORA

In our opinion, there are two fundamental qualities that a specification language should have: the language must be easy to comprehend (a specification has more readers than writers) and the users must like it.

Therefore, we experimentally validated the ADORA language with respect to these two qualities. We conducted an experiment with the following goals.

- Determine the comprehensibility of an ADORA specification both on its own and in comparison with an equivalent specification written in UML – today’s standard modeling language – from the viewpoint of a reader of the specification.
- Determine the acceptance of the fundamental concepts of ADORA (using abstract objects, hierarchical decomposition, integrated model...) both on its own and in comparison with UML from the viewpoint of a reader/writer of models.

4.1 Setup of the Experiment

In order to measure these goals, we set up the following experiment [2]. We wrote a partial specification of a distributed ticketing system both in ADORA and in UML. The system consists of geographically distributed vending stations where users can buy tickets for events (concerts, musicals...) that are being offered on several event servers. Vending stations and event servers shall be connected by an existing network that needs not to be specified.

Then we prepared a questionnaire consisting of two parts. In the first part, the “objective” one, we aimed at measuring the comprehensibility of an ADORA model. We created 30 questions about the contents of the specification, for example “Can a user at a point of sale terminal purchase an arbitrary number of tickets for an event in a single transaction?” 25 questions were yes/no questions; the rest were open questions. For every question, we additionally asked, whether the answering person was sure or unsure about her or his answer and how difficult it was to answer the question.

In the second part, the “subjective” one, we tested the acceptance of ADORA vs. UML. We asked 14 questions about the personal opinion of the answering person concerning distinctive features of both ADORA and UML, for example “Does it make sense to use an integrated model (like ADORA) for describing all aspects of a system”?

We ran the experiment with fifteen graduate and Ph.D. students in Computer Science who were not members of our research group. The participants were first given an introduction both to ADORA and to UML. Then we divided the participants into two groups. The members of group A answered the objective part of the questionnaire first for the ADORA specification and then for the UML specification; group B members did it vice-versa. Finally, both groups answered the subjective part of the questionnaire. In order to avoid answers being biased towards ADORA, we ensured the anonymity of the filled questionnaires.

Two participants did not finish the experiment; another person's answers could not be scored because his answers revealed insufficient base knowledge of object technology. So we finally had twelve complete sets of answers.

4.2 Some Results

Due to space limitations, we only can present some key results here. The complete results are given in [2]. As the differences between groups A and B are marginal, we consolidate the results for both groups in the results given below.




Fig. 9. Comprehensibility of models. Right and wrong answers to the questions in the objective part of the questionnaire for ADORA vs. UML models. The graphics also shows how certain the participants were about their answers and how they rated the difficulty of answering.

Fig. 9 shows the overall results of the first part of the questionnaire. For each model, we had a total of 360 answers (30 questions times 12 participants). For every answer, we determined whether the answer was objectively right or wrong. The answers were further subdivided into those where the answering person was sure about her or his answer and those where she or he was not. The subdivision of the columns indicates how difficult it was to answer the questions in the participants' opinion. (For example, about 79% of the questions about the ADORA model were answered correctly and the participants were sure about their answer. For about half of these answers, the participants judged the answer to be easy to give.)

Despite the fact that the number of participants was fairly small, these results strongly support the comprehensibility hypothesis and also show a clear trend that an

ADORA specification is easier to comprehend than an UML specification. Moreover, we have two important results that are statistically significant at a level of 0.5% [2]:

- The percentage of correctly answered questions is higher for the ADORA model than for the UML model. That means, reading the ADORA model is less error-prone than reading the UML model.
- When answering a question correctly, the readers of the ADORA model are more confident of themselves than the readers of the UML model.

Table 1 summarizes the results of the subjective part of the questionnaire. Again, the results strongly support our hypothesis that users like the fundamental concepts of ADORA and that they prefer them to those of UML.

Table 1. Acceptance of distinct features; ADORA vs. UML

Statement		strongly agree	mostly agree	mostly disagree	strongly disagree
The specification gives the reader a precise idea about the system components and relationships	ADORA	23%	62%	8%	8%
	UML	8%	46%	31%	15%
The structure of the system can be determined easily	ADORA	54%	31%	8%	8%
	UML	8%	38%	23%	31%
The specification is an appropriate basis for design and implementation	ADORA	25%	75%	0%	0%
	UML	0%	50%	33%	17%
Using an integrated model (ADORA) makes sense		42%	25%	33%	0%
Using a set of loosely coupled diagrams (UML) makes sense		8%	17%	67%	8%
Hierarchical decomposition eases description of large systems		15%	69%	15%	0%
	ADORA eases focusing on parts without losing context	38%	46%	15%	0%
Decomposition in ADORA eases finding information		46%	38%	15%	0%
Integrating information from different diagrams is easy in UML		15%	15%	46%	23%
Specifying objects with their roles and context is adequate		31%	54%	15%	0%
Describing classes is sufficient		0%	15%	62%	23%

Even if we subtract some potential bias (maybe some of the participating students did not want to hurt us), we can conclude from this experiment that the ADORA language is clearly a step into the right direction.

5 Yet Another Language? ADORA vs. UML

The goal of the ADORA project is not to bless mankind with another fancy modeling language. When UML became a standard, we of course investigated the option of making ADORA a variant of UML. The reason why we didn't is because ADORA and UML differ too much in their basic concepts (Table 2). The most fundamental difference is the concept of an integrated, hierarchically decomposable model in ADORA vs. a flat, mostly non-decomposable collection of models in UML. Using packages for an ADORA-like decomposition fails, because UML packages are mere containers. Using stereotypes for integrating ADORA into UML would be an abuse of this concept, as the resulting language would no longer behave like UML. A real integration of ADORA-concepts into UML would require major changes in the UML metamodel [9].

Table 2. Comparison of basic concepts of ADORA vs. UML

ADORA	UML
Specification is based on a model of abstract objects, types are supplementary	Specification is based on a class model, object models are partial and supplementary
Specifies all aspects in one integrated model; separation of concerns achieved by decomposition and views	Uses different models for each aspect. Separates concerns by having a loosely coupled collection of models
Hierarchical decomposition of objects is the principal means for structuring and comprehending a specification	Class and object models are flat. Only packages can be decomposed hierarchically
Scenarios are tightly integrated into the specification; they can be structured and decomposed systematically	Use cases (=type scenarios) are loosely integrated with class and object models. Structuring capabilities are weak, decomposition is not possible.
Precise rules for consistency between aspect views	Nearly no consistency rules between aspect models
Conceptual visualization eases orientation and navigation in the specification and improves comprehensibility	UML tools provide traditional scrolling and explosive zooming only

6 Conclusions

Summary. We have presented ADORA, an approach to object-oriented modeling that is based on object modeling and hierarchical decomposition, using an integrated model. The ADORA language is intended to be used for requirements specifications and high-level, logical views of software architectures.

Code generation. ADORA is not a visual programming language. Therefore, we have not done any work towards code generation up to now. However, in principle the generation of prototypes from an ADORA model is possible. ADORA has both the structure and the language elements that are required for this task.

State of work. We have finished a first definition of the ADORA language in 1999 [13]. In the meantime we have evolved some language concepts and have conducted an experimental validation. The ADORA tool is still in the proof-of-concept phase. We have a prototype demonstrating that the zooming algorithm, which is the basis of our visualization concept, works.

Future plans. The work on ADORA goes on. In the next years, we will develop a real tool prototype and investigate the use of ADORA for partial and incrementally evolving specifications. Parallel to that, we want to apply ADORA in projects and evolve the language according to the experience gained.

References

1. Berner, S., Joos, S., Glinz, M., Arnold, M. (1998). A Visualization Concept for Hierarchical Object Models. *Proceedings 13th IEEE International Conference on Automated Software Engineering (ASE-98)*. 225-228.
2. Berner, S., Schett, N., Xia, Y., Glinz, M. (1999). *An Experimental Validation of the ADORA Language*. Technical Report 1999.05, University of Zurich.
3. Booch, G. (1994). *Object-Oriented Analysis and Design with Applications*, 2nd ed. Redwood City, Ca.: Benjamin/Cummings.

4. Carroll, J.M. (ed.)(1995). *Scenario-Based Design*. New York: John Wiley & Sons.
5. Coad, P., Yourdon E. (1991). *Object-Oriented Analysis*. Englewood Cliffs, N. J.: Prentice Hall.
6. Firesmith, D., Henderson-Sellers, B. H., Graham, I., Page-Jones, M. (1998). *Open Modeling Language (OML) – Reference Manual*. SIGS reference library series. Cambridge: Cambridge University Press.
7. Glinz, M. (1993). Hierarchische Verhaltensbeschreibung in objektorientierten Systemmodellen – eine Grundlage für modellbasiertes Prototyping. [Hierarchical Description of Behavior in Object-Oriented System Models – A Foundation for Model-Based Prototyping (in German)] In Züllighoven, H. et al. (eds.): *Requirements Engineering '93: Prototyping*. Stuttgart: Teubner. 175-192.
8. Glinz, M. (1995). An Integrated Formal Model of Scenarios Based on Statecharts. In Schäfer, W. and Botella, P. (eds.): *Software Engineering – ESEC'95*. Berlin: Springer. 254-271.
9. Glinz, M. (2000). Problems and Deficiencies of UML as a Requirements Specification Language. *Proceedings of the Tenth International Workshop on Software Specification and Design*. San Diego. 11-22.
10. Harel, D. (1987). Statecharts: A Visual Formalism for Complex Systems. *Sci. Computer Program*. **8** (1987). 231-274.
11. Jacobson, I., Christerson, M., Jonsson, P., Övergaard, G. (1992). *Object-Oriented Software Engineering – A Use Case Driven Approach*. Reading, Mass.: Addison-Wesley.
12. Joos, S., Berner, S., Arnold, M., Glinz, M. (1997). Hierarchische Zerlegung in objektorientierten Spezifikationsmodellen [Hierarchical Decomposition in Object-Oriented Specification Models (in German)]. *Softwaretechnik-Trends*, **17**, 1 (Feb. 1997), 29-37.
13. Joos, S. (1999). *ADORA-L – Eine Modellierungssprache zur Spezifikation von Software-Anforderungen* [ADORA-L – A modeling language for specifying software requirements. (in German)]. PhD Thesis, University of Zurich.
14. Leveson, N.G., Heimdahl, M.P.E., Reese, J.D. (1999). Designing Specification Languages for Process Control Systems: Lessons Learned and Steps to the Future. In Nierstrasz, O. and Lemoine, M. (eds.): *Software Engineering – ESEC/FSE'99*. Berlin: Springer. 127-145.
15. Rumbaugh, J., Blaha, M., Premerlani, W., Eddy, F., Lorensen, W. (1991). *Object-Oriented Modeling and Design*. Englewood Cliffs, N. J.: Prentice Hall.
16. Rumbaugh, J., Jacobson, I., Booch, G. (1999). *The Unified Modeling Language Reference Manual*. Reading, Mass.: Addison-Wesley.
17. Schett, N. (1998). *Konzeption und Realisierung einer Notation zur Formulierung von Integritätsbedingungen für ADORA-Modelle*. [A notation for integrity constraints in ADORA models – Concept and implementation (in German)]. Diploma Thesis, Univ. of Zurich.
18. Selic, B., Gullekson, G., Ward, P. T. (1994). *Real-Time Object-Oriented Modeling*. New York: John Wiley & Sons.
19. Wirfs-Brock, R., Wilkerson, B., Wiener, L. (1993). *Designing Object-Oriented Software*. Englewood Cliffs, N. J.: Prentice Hall.