

Adaptability Aspects: An Architectural Pattern for Structuring Adaptive Applications with Aspects

Ayla Dantas* and Paulo Borba†
Informatics Center – Federal University of Pernambuco
Po Box 7851 - 50.732-970 Recife, PE - Brazil
{add,phmb}@cin.ufpe.br

Abstract

*This paper presents an architectural pattern for structuring adaptive applications using aspect-oriented programming in order to obtain separation of concerns. It is composed of known and novel patterns organized so as to provide good maintainability and modularity.*¹

1 Intent

This architectural pattern is intended to show how to use aspects [14] in order to better structure adaptive applications, which are able to change their behavior in response to context changes [11], such as the device new localization or its resources state.

2 Context

Adaptability has become a common requirement [9] and its implementation usually affects many parts of the code. Most implementations of this requirement lead to tangled code, mixing different concerns such as application business rules, GUI code, and adaptive behavior implementation. It is sometimes hard to include the adaptability concern in new and existing applications in easily maintainable way. Besides that, the mechanisms for accessing contextual information, such as environment sensors or analyzers, change frequently, which usually demands regular modifications to the application that can be hard to do. This is a problem because contextual information obtained by new sources may lead to different application behaviors. For example, we can have a device that becomes able of accessing its temperature, and according to it changes its processing mode.

Copyright ©2003, Ayla Dantas and Paulo Borba. Permission is granted to copy for the Sugarloaf-PLoP 2003 Conference. All other rights are reserved.

*Supported by CNPq.

†Partially supported by CNPq, grant 521994/96–9.

¹We followed the POSA (Pattern-Oriented Software Architecture) structure to present our pattern, including an Intent section and distributing the contents of sections to be called *Variants* and *Resolved Example* throughout the pattern.

3 Problem

Avoid lack of flexibility and code tangling mixing adaptive functionality implementation with different concerns such as application business rules and GUI code.

4 Forces

In order to solve the problem, *Adaptability Aspects* balances the following forces:

- Separate adaptability concerns from other concerns, in order to improve reuse and maintainability.
- The adaptability functionality might be either plugged in/out and also turned on/off, because the ways to access the context and the behaviors triggered by environment state may change a lot.
- Not all the developers need to know a particular Aspect-Oriented Programming (AOP) language, because each team may be focused on a different pattern module, that does not necessarily need AOP.
- The application should be easy to maintain.
- The application can be implemented in any platform, from embedded devices, such as cellular phones, to enterprise applications.
- The kind of contextual information might change and this should not cause a significant impact on the system.

5 Solution

Use aspects [14] to make the application adaptive in a modular and not invasive way. With Aspect-Oriented Programming we can cleanly capture some implementation aspects that affect many parts of a system, as it happens with adaptability, providing appropriate isolation, composition and reuse of the code used to implement those aspects [8].

By using aspects in *Adaptability Aspects* we define how the behavior of the application functionalities should be changed in order to support adaptability. These aspects are able to crosscut some application execution points and change their normal execution. They should interact with a module for monitoring the context in order to identify when a change must be triggered and they use auxiliary classes to perform the application changes. For a better flexibility, the auxiliary classes should interact with a module responsible for obtaining dynamic data specifying how the application should adapt in a specific situation. Auxiliary classes are used to avoid requiring all developers to know an AOP language and to avoid code tangling in the adaptive behavior implementation.

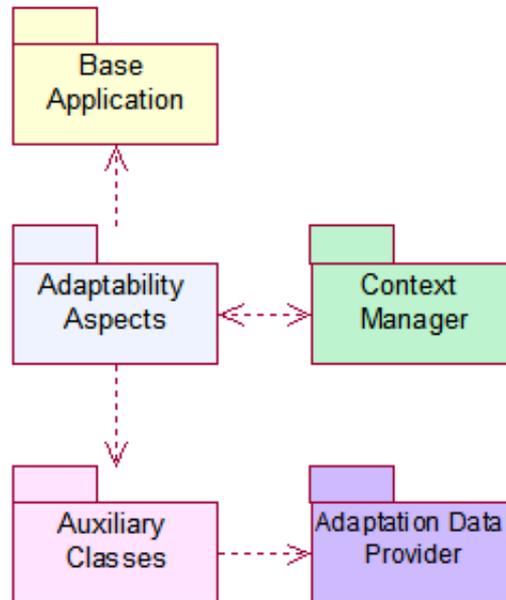


Figure 1: *Adaptability Aspects* elements.

6 Structure

The *Adaptability Aspects* architectural pattern presents five elements or modules:

- **Base Application:** The core application functionalities, such as business and GUI code, and possibly persistence and distribution code, but no adaptability code.
- **Adaptability Aspects:** The aspects implementing the adaptability concern. They specify how the behavior of the base application functionalities should be changed to adapt to contextual changes. This element delegates several tasks to the auxiliary classes.
- **Auxiliary Classes:** Classes used by the aspects to provide the adaptive behavior. Its isolation from the aspect is intended to improve reuse. They communicate with the adaptation data provider module in order to obtain dynamic data for the adaptation. Besides that, for developing the auxiliary classes, the developers of this module do not necessarily need to know an AOP language. The Adaptability Aspects developer or the system architect may simply specify the interfaces of these classes and what they should do. Then, from these specifications, these classes can be built and have their methods invoked by the aspects.
- **Context Manager:** Module responsible for analyzing context changes and triggering adaptive actions implemented by the aspects. It can also be called by the aspects to obtain information about the context. Its implementation can be based on a variation of the Observer pattern [4], or on its implementation with aspects [6]. In this way, new mechanisms for accessing the context can be easily supported without significant impact on the application (see the Example section).

- **Adaptation Data Provider:** Classes responsible for providing data for dynamic adaptations according to context changes. This means that the same context change can lead to different behaviors in different moments according to the data provided by this module. These classes can be organized as an Adaptive Object-Model (AOM) [16].

These elements and their inter-relation are shown in Figure 1. They are represented there using the UML package notation. Each package represents a logical part of the code, but each of these parts can be implemented using several programming language-specific packages. The arrows represent the dependency between the packages. It is important to notice that the relation between *Context Manager* and *Adaptability Aspects* is double-way but this does not imply in internal packages double-way relations. Sometimes a *Context Manager* package uses an *Adaptability Aspects* package to notify about context changes or an *Adaptability Aspects* package may use a *Context Manager* to request some information about the context in a certain execution point. Another important observation is that we have a different kind of dependency between *Base Application* and *Adaptability Aspects* because the latter is capable of crosscutting the former.

7 Dynamics

The following scenarios depict the dynamic behavior of the *Adaptability Aspects* pattern.

Scenario I, which is illustrated by Figure 2, presents the application behavior when the aspects request information about the context at specific points in the execution flow and the application changes its behavior at those points according to the response obtained:

- The application starts executing one of its core functions.
- The adaptability aspects detect points in the execution flow (join points) of the core function where an adaptation might be performed (secure points). This is implemented using AOP constructs such as pointcuts, which are a way of identifying join points. Then, these aspects lead to behavior changes in the application before, after or around those points using advice.
- Those aspects then query the context manager to verify whether any adaptation should be performed, which will depend on the environment state.
- The actions implementing the adaptation are then delegated to the auxiliary classes. These actions are performed before, after, or around the join points.
- In order to access dynamic data specifying how an adaptation should be performed, the auxiliary classes access the objects of the adaptation data provider, which can be based on the Metadata and Active Object-Model pattern language [3], also called Adaptive Object-Model Architecture [15].
- The application is then changed according to this dynamic data.

The sequence diagram shown in Figure 2 shows the interaction between the pattern elements. Instead of representing objects, each diagram box represents a collection of objects from the correspondent pattern module.

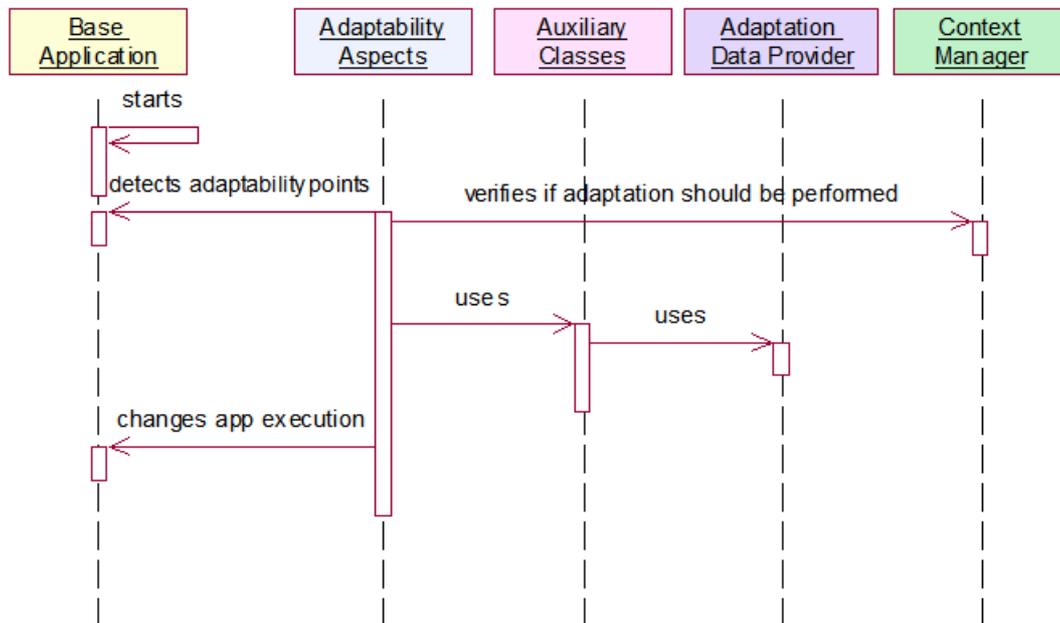


Figure 2: Dynamics of *Adaptability Aspects* pattern (scenario I)

Scenario II, illustrated by Figure 3, shows how the elements of the *Adaptability Aspects* pattern interact when a context change triggers an adaptation on the application behavior.

- The application starts execution.
- The context manager begins to monitor the context continuously.
- When a registered environment change is detected, the adaptability aspects are notified.
- The adaptability aspects detect a point in the execution flow and then use auxiliary classes to perform the changes in the application.
- The auxiliary classes use the adaptation data provider in order to access dynamic information on how the adaptation should be performed.
- The aspects change the application behavior.

When the adaptability aspects are notified about a context change, the action to be performed can execute either immediately, as for example, an invocation to the garbage collector or not, as it is more usual. In the latter case, a field representing the adaptation state can be directly set, but the adaptive behavior is introduced on the base application just when a specific join point is reached.

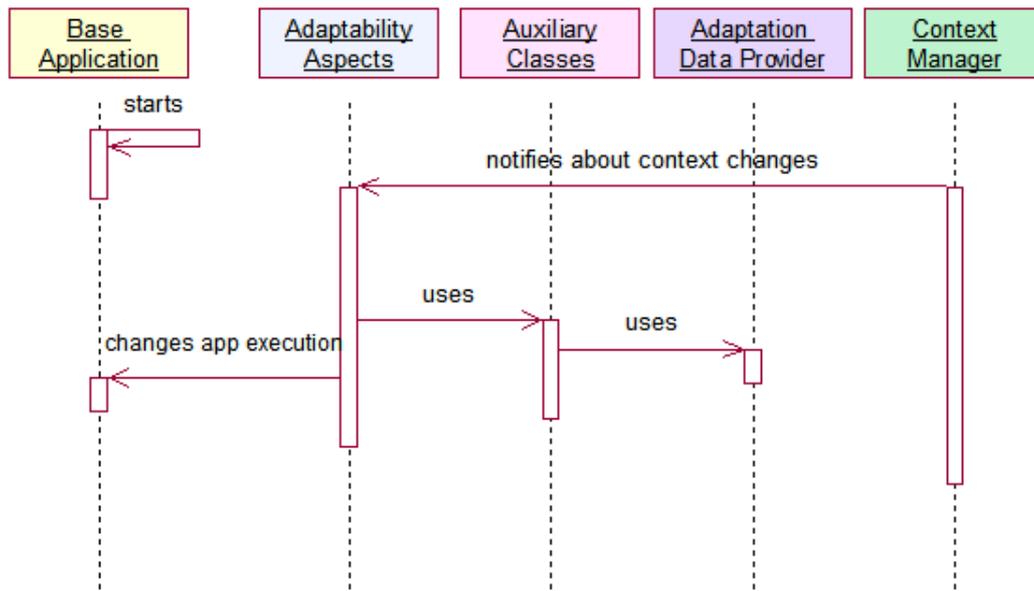


Figure 3: Dynamics of the *Adaptability Aspects* pattern (scenario II)

8 Consequences

The *Adaptability Aspects* pattern provides the following benefits:

- *Modularity.* The base application is isolated from the adaptability aspects and from the classes that actually execute actions at the identified adaptation points.
- *Reuse.* The use of auxiliary classes improves reuse, since these classes can be used by many aspects. Besides that, as will be seen on the implementation section, the Adaptation Data Provider and Context Manager modules should also be organized in order to improve reuse by other adaptive applications.
- *Extensibility.* As the aspects code is isolated from the core and auxiliary classes code, it becomes easier to maintain each part of the application and also to extend the system. The context manager should also be internally organized to promote extensibility, making it simpler to achieve and minimizing the impacts caused by the inclusion of a new context element. Nevertheless, the use of AOP in some modules must be done carefully and the use of visual tools is strongly recommended, because AOP languages are considerably powerful, and a change in one part of its code can affect the entire system.
- *Platform Independence.* Its general structure can be applied to a wide range of systems, from embedded systems (such as the example we will present shortly) to enterprise applications. This follows from it not imposing a burden on efficiency and not requiring a significant amount of resources.

- *Dynamic changes.* With the adaptation data provider module, dynamic changes in the application behavior can be performed, and those changes do not need to be expressed by new code; they may be expressed in metadata (XML files, for example).

The *Adaptability Aspects* architectural pattern also has some liabilities:

- *Code size.* Implementations of the pattern are usually bigger than alternative non-modular solutions that tangle adaptation code with base application code, as we can observe from design and based on our implementations. The implementations proposed on the Implementation section for this architectural pattern can also increase the code size if more flexibility, dynamicity, and reuse are required. This liability can be a problem for embedded systems, in which resources are restricted, but for enterprise applications it would not be an issue.
- *Efficiency.* Due to the number of elements required by the architecture, the proposed pattern imposes a burden on efficiency when compared to other non-modular solutions. Nevertheless, such burden would be a minor problem for an enterprise architecture. Besides that, some efficiency problems may be related to the used AOP language.
- *Dynamic loading.* An important requirement related to adaptability nowadays is dynamicity (the extent to which the adaptation should be dynamic). Our architectural pattern provides dynamic changes using the adaptation data provider element, which presents new application behaviors in metadata. However, it may not be possible to dynamically load code in order to include new kinds of adaptation mechanisms (ways of triggering new adaptations) without stopping the application, which will depend on the application platform and on the AOP language being used. AspectJ [7], an AOP language in widespread use, does not allow dynamic loading of new aspects in its latest stable version.
- *New programming paradigm.* As we have at least one module in the pattern that uses aspects, some developers of the team need to know Aspect-Oriented Programming. Learning this new programming paradigm may lead to more development time, initially. However, for other modules, or even for some parts of a module, knowing a particular AOP language is not required.

9 Implementation

The following guidelines help implementing an adaptive system using the *Adaptability Aspects* pattern.

1. *Develop the base application.* The application is developed with its main functionalities, preferably organized in a way to help its evolution. If the application has already been developed, it is occasionally suggested to do some refactoring in order to provide internal modularity to this module.
2. *Develop the context management module.* This module must present ways of detecting which elements from the context should change, which elements should be

informed when these changes occur, and which actions should take place at those moments. We propose the implementation of this module using the *Observer* pattern, and specially its implementation using aspects, proposed in another work [6], and adapted to the observation of the context, which increases the flexibility of this module in relation to new environment elements to be observed. This is illustrated by Figure 4. Although this figure illustrates the Context Manager module, the `SpecificAdaptabilityAspect` is part of the Adaptability Aspects part and is shown here just to illustrate the interaction between these pattern elements. The `SpecificAdaptabilityAspect` and `SpecificAdaptationProtocol` aspects and the `SpecificContextElementVerifier` class of this figure are used to represent the elements used to manage a specific context change. For example, if we want a different application behavior, such as the changing of languages used for translation in a dictionary, when the device localization changes, we would replace these classes by the following ones respectively: `DictionaryCustomization`, `LanguageAdaptationProtocol` and `LocalizationVerifier`. Figure 7, in the next section, better illustrates this example and the inter-relationship among the pattern elements in this situation.

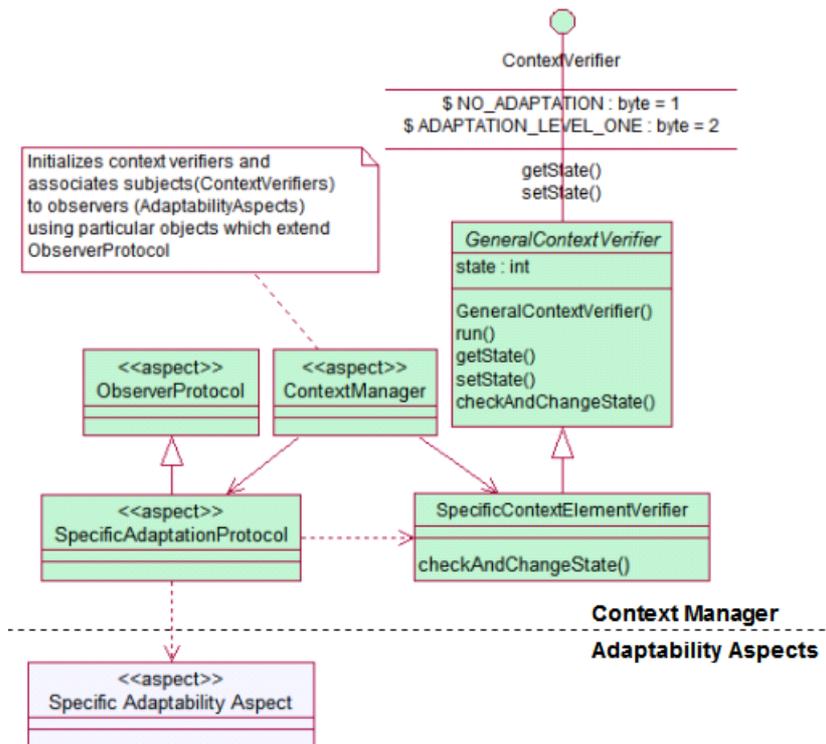


Figure 4: Context Manager Possible Implementation

3. *Develop the adaptation data provider module.* Dynamic adaptation is generally desirable. It is not good practice to fix the adaptations in code. One option is to have these changes in metadata such as XML files that can be obtained locally or remotely (the Bridge [4] pattern is indicated for providing this flexibility). Such an idea is adopted by the Adaptive Object-Model architecture [16]. The use of AOM

is not mandatory for the development of this module, but presents the advantage that, once implemented, one of its parts can be reused by many systems; only the interpreter of this AOM remains to be implemented according to the wanted changes, which is a task that can be performed by auxiliary classes. An AOM partial implementation that can be reused by many kinds of adaptive applications is shown in Figure 5. The application elements can be represented as **EntityType** elements, which present **Entities** with their **Properties** or **StrategyObjects** to represent different system behaviors. The interpreter of the AOM will translate these kinds of elements into application properties or algorithms. The auxiliary classes used by the adaptability aspects will have access to the adaptation data provider module by the **AppAOMManager** class, using its **getEntityType** method. An alternative implementation for the adaptation data provider module may explore the idea presented in another work [2], using an application-specific object responsible for obtaining dynamic data for adaptations and making them available through its operations. However, with this approach, besides losing in reuse, for each new kind of data, we would have to modify this object hierarchy, and the number of operations could increase a lot.

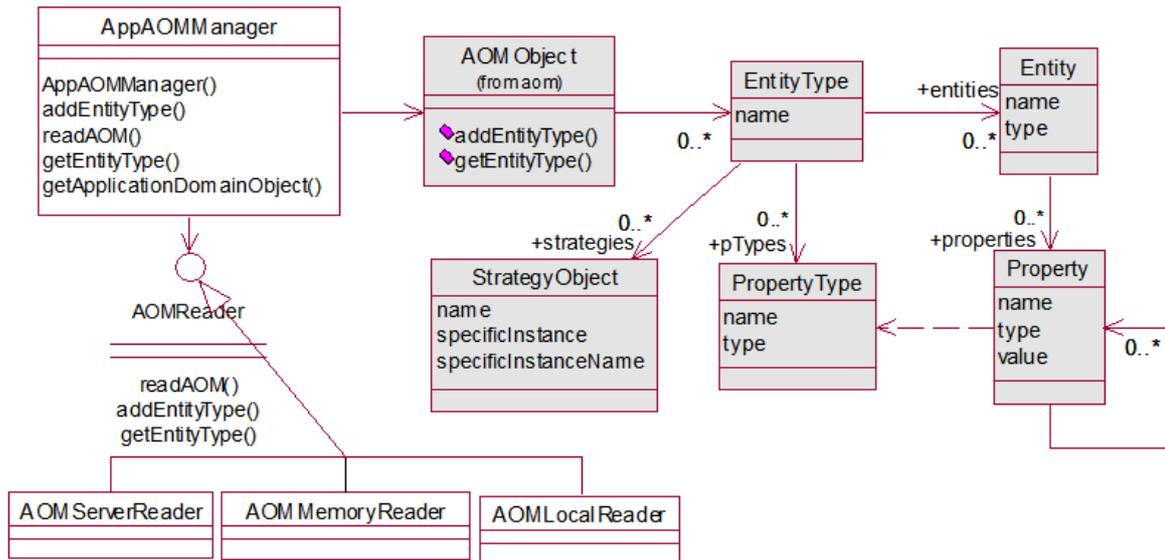


Figure 5: Adaptation Data Provider using AOM

4. *Identify which kind of adaptations will be included.* It is necessary to analyze what will really constitute an adaptation (an application change caused by the environment) and which context changes would lead to those adaptations. Example adaptations are changes in the application language due to the detection of a device location modification, memory management, changes and inclusion of screens in the application due to user inputs or server responses, for example.
5. *Transform each kind of adaptation into an aspect.* Use AOP constructions such as pointcuts and advice to access the application components that can change and the execution points (join points) where the adaptations should be performed. These

aspects should contain the least amount of non-aspect code as possible, delegating necessary actions to auxiliary classes that will be developed in a later step. Besides that, these classes should be able to access the context manager module in order to verify whether an adaptation should be done. Each aspect can also be a `SpecificAdaptabilityAspect` in the situation illustrated by Figure 4.

6. *Develop the auxiliary classes.* These classes will actually present the code to be executed at the points selected in the aspects module. Their design should promote reuse by many aspects. In order to perform dynamic adaptations, this module should communicate with the adaptation data provider module. One suggestion is to include in these classes interpreters of AOM for the application being developed.

10 Example

To exemplify the *Adaptability Aspects* pattern, and specially its second scenario, we will now consider a dictionary application developed in Java [5], for the J2ME (Java 2 Micro Edition) platform. This platform accommodates consumer electronics and embedded devices [12]. The AOP language used is AspectJ [7], a general purpose aspect-oriented extension to Java in widespread use.

Our **Base Application** will be composed by this dictionary, which is a simple MIDP (Mobile Information Device Profile)-based application (also known as MIDlet [12]), without any adaptations. Its basic functionality is to translate a given word from English into Portuguese. It follows the MVC [1] architectural pattern and its main classes are illustrated by Figure 6: the controller part is being represented by the `DictionaryMIDlet` and the `DictionaryController`; the view part is represented by the `DictionaryScreen`; the model is represented by the `InputSearchData` class, which stores some information about the dictionary, such as the word being translated and the source and destination languages of the translation.

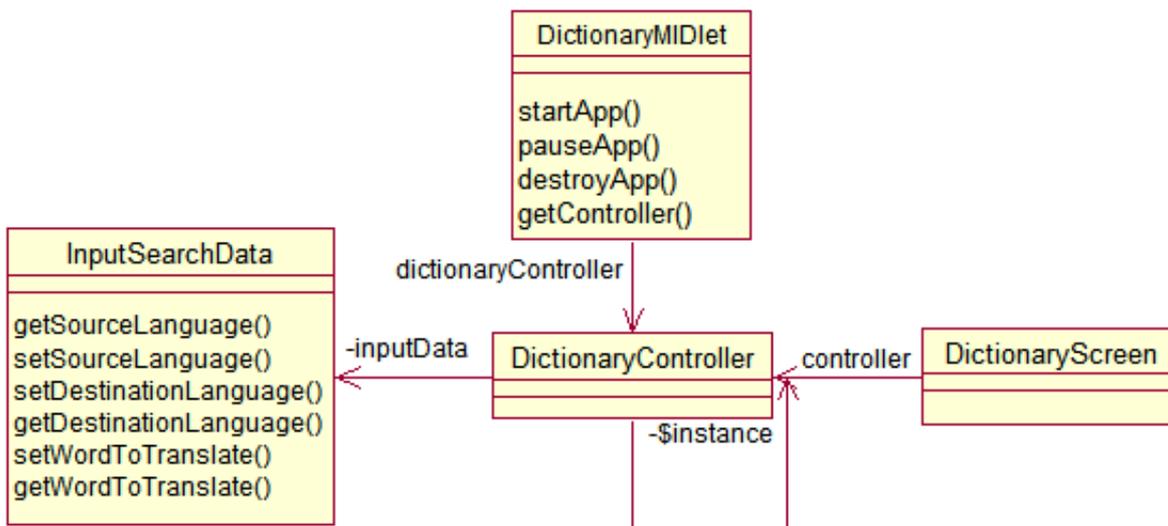


Figure 6: A **Base Application** example

In order to make the example easier to understand, the adaptation we will include will be the change of the application translation language when the device location changes. The main classes and aspects used for this adaptation are illustrated by Figure 7, which is a UML diagram that uses the `<< aspect >>` stereotype to identify aspects. Other stereotypes are also used in some dependency relationships to represent classes whose behavior is monitored or changed by an aspect (`<< affects >>` stereotype), or those which are just used as auxiliary classes (`<< uses >>` stereotype) by an aspect.

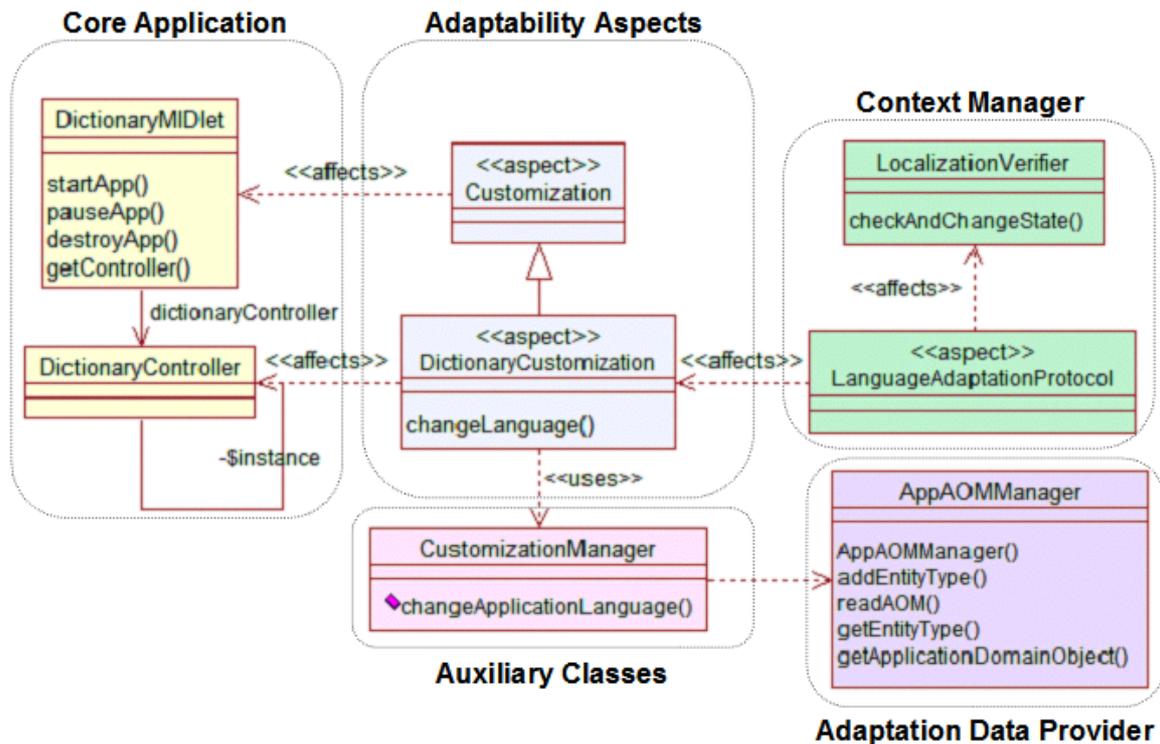


Figure 7: An **Adaptability Aspects** pattern example

The aspects which will compose the **Adaptability Aspects** element of the pattern are organized in a small framework, composed of two aspects: the `Customization` aspect (which can be reused by any MIDP-based J2ME application) and the `DictionaryCustomization` aspect (which is application specific and extends the former). They are responsible for changing application properties, using a captured `MIDlet` instance at the moment the `startApp` method is executed. The following code from the `Customization` aspect illustrates that.

```
pointcut MIDletStart(MIDlet midlet):
    execution(void startApp()) && target(midlet);
before(MIDlet midlet): MIDletStart(midlet) {
    adaptBefore(midlet);
}
```

The **Customization** aspect is general because any MIDP-based J2ME application should present a class implementing `MIDlet` with a `startApp` method. Its `before` advice invokes the `adaptBefore` abstract method, which is defined in the `DictionaryCustomization` aspect. At its invocation, this aspect stores the `MIDlet` instance, which makes it able of changing application properties in its methods. The interaction between these aspects and the application classes is illustrated by Figure 7.

The **Context Manager** module will follow the structure shown in Figure 4. The code for the `ObserverProtocol` is shown elsewhere [6]. Our `SpecificAdaptationProtocol` aspect will be the `LanguageAdaptationProtocol`, which is shown below:

```

1: public aspect LanguageAdaptationProtocol extends ObserverProtocol {
2:   declare parents: LocalizationVerifier implements Subject;
3:   declare parents: DictionaryCustomization implements
4:     Observer;
5:   protected pointcut subjectChange(Subject s):
6:     call(void LocalizationVerifier.setState(int)) && target(s);
7:   protected void updateObserver(Subject s, Observer o) {
8:     LocalizationVerifier lv = (LocalizationVerifier) s;
9:     if (lv.getState() != ContextVerifier.NO_ADAPTATION) {
10:      ((DictionaryCustomization) o).changeLanguage();
11:      ((LocalizationVerifier)
12:        s).setState(ContextVerifier.NO_ADAPTATION);
13:    }
14:  }
15:}

```

In this aspect we identify the `Subject` (`LocalizationVerifier`) and the `Observer` (`DictionaryCustomization` aspect) of the pattern using AspectJ introduction (lines 2 and 3). Then, we define the pointcut `subjectChange`, inherited from `ObserverProtocol`, identifying the execution points that characterize the `Subject` change (lines 5 and 6), which is the call of the `setState` method on a `LocalizationVerifier` object in this example. Finally, we define the also inherited method, `updateObserver`, indicating what should happen to the `Subject` and the `Observer` when there is a change (lines 7-14). In this case, a method from the `DictionaryCustomization` aspect is invoked to change the application language. We will talk about this method later, while describing our **Auxiliary Classes** module use.

The `GeneralContextVerifier` from Figure 4 extends the `Thread` class and invokes an abstract method called `checkAndChangeState` at regular intervals. This class can be reused by other applications. The `LocalizationContextVerifier` extends the `GeneralContextVerifier` and should only specify its abstract method, verifying the device localization as the following code illustrates.

```

private LocalizationObject lo = new LocalizationObject();
public void checkAndChangeState() {
  if (lo.hasChanged())
    this.setState(ContextVerifier.ADAPTATION_LEVEL_ONE);
}

```

We use here an abstraction of a `LocalizationObject`, which is responsible for verifying if the device localization has changed. If this happens, the `setState` method is called, the `LanguageAdaptationProtocol` aspect identifies that the `Observer` should be notified and executes the `updateObserver` method, which will call the `changeApplicationLanguage` method from the `DictionaryCustomization` aspect.

As Figure 4 illustrates, the `ContextManager` initializes the context verifiers and associates subjects to observers. The following code extracted from it shows how it works for this example.

```
public aspect ContextManager {
    pointcut MIDletStart(MIDlet midlet): execution(void startApp()) &&
        target(midlet);
    after(MIDlet midlet): MIDletStart(midlet){
        LocalizationVerifier lv = new LocalizationVerifier();
        lv.start();
        LanguageAdaptationProtocol.aspectOf().addObserver(lv,
            DictionaryCustomization.aspectOf());
    }
    ...
}
```

In order to avoid non-aspect code in aspects, the `changeLanguage` method from the `DictionaryCustomization` aspect delegates the application language change to an auxiliary class, the `CustomizationManager`, which will be part of the **Auxiliary Classes** element of the *Adaptability Aspects* pattern. Nevertheless, this element is not mandatory and another possible technique would be to use non-aspect code in methods of the aspect and not in the advice body when extensive code is necessary to provide the change. However, to illustrate the **Auxiliary Classes** module use and its interaction with **Adaptation Data Provider** module (see Figure 5), we will show one of the `CustomizationManager` methods implementation, the `changeApplicationLanguage`:

```
public void changeApplicationLanguage(){
    EntityType et = aomManager.getEntityType("DictionaryProperties");
    Entity e = et.getEntity("InputSearchData");
    Property sourceLang = e.getProperty("sourceLanguage");
    Property destLang = e.getProperty("destinationLanguage");
    InputSearchData isd = this.midlet.getController().getInputData();
    isd.setDestinationLanguage(destLang.getValue().toString());
    isd.setSourceLanguage(sourceLang.getValue().toString());
}
```

In this code we can notice the relation of the **Auxiliary Classes** component with the **Adaptation Data Provider** module. Here, this module uses the *Adaptive Object-Model* architectural style. As we can see, the application properties and behavior, described in metadata (in our case an XML file), are translated into the object structure illustrated by Figure 5. An UML object diagram of a possible configuration for this code is illustrated by Figure 8. The `changeApplicationLanguage` method then interprets this object model

and changes the application using the MIDlet object and modifying the `InputSearchData` instance. This is just a simple example of AOM's use. By using it, better exploring the patterns it is related to, we can do much more things that would lead to more flexibility.

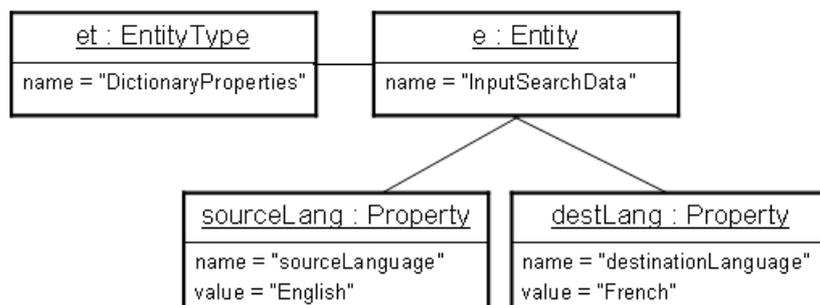


Figure 8: Entity configuration example

11 Known Uses

The *Adaptability Aspects* architectural pattern has been used in two experiments we have developed with J2ME: a Dictionary application for a cellular phone and an album for pictures, also for small devices. Besides that, some parts of the patterns have already been referred to. An example of this is a work [10] presenting some experiments using aspects. One of the conclusions is that it is an interesting architecture to separate an AOP system in three parts: the base code, the aspects and the auxiliary code. In this structure, the aspects part functions as glue between the other two. Another example is the implementation of the Observer aspect using aspects proposed by Hanneman [6] which we modify here in the **Context Manager** module.

12 See Also

- The *Reflection* architectural pattern [1], which provides a mechanism for changing structure and behavior of software systems dynamically, is related to our pattern. It is intended to make applications adaptable, that means, able to easily evolve due to requirement changes, which is also one of our requirements. Although it provides a lot of flexibility, it seems to increase the complexity of the system more than our solution and even to be less efficient.
- *Adaptive Object-Model* systems represent their attributes, classes, and relationships as metadata [15]. AOMs can be used to represent all the system structure or just what changes in the system. In our pattern, we took the latter approach and explored AOP as a non-invasive way to include AOMs. We have used AOMs in the **Adaptation Data Provider** module as the provider of the adaptive behavior data.
- The *PADA (Pattern for Distribution Aspects)* [13], even dealing with Distribution, is also related to this work as it provides a solution to the lack of modularity and

maintainability using AOP.

- The *Observer* pattern [4] is very useful in the implementation of the **Context Manager** component (see Figure 4), which actually characterizes the adaptive behavior. Its implementation with aspects [6] is interesting as a dynamic way (using pointcuts) of identifying subject changes.
- Another work [2] describes some practices incorporated to this pattern and presents the implementation of some adaptive concerns into J2ME applications using AspectJ, comparing them to pure Object-Oriented solutions using GoF patterns. The implementation shown there does not exactly follow this pattern and does not use Adaptive Object-Models. Instead of that, it uses a less flexible approach for obtaining dynamic data for adaptations.

13 Acknowledgments

We would like to give special thanks to Paulo Masiero, our shepherd, for his important comments and suggestions, helping us to improve our pattern. Thanks to Rossana Andrade for her work as mediator and her attention, and also to Vander Alves, for his valuable suggestions. Besides them, we also want to thank the conference participants for the suggestions made at the Writers' Workshop, and specially Joseph Yoder, for his comments and ideas even after the SugarLoafPloP.

References

- [1] Frank Buschmann, Regine Meunier, Hans Rohnert, Peter Sommerlad, and Michael Stal. *A System of Patterns: Pattern-Oriented Software Architecture*. John Wiley & Sons, 1996.
- [2] Ayla Dantas and Paulo Borba. Developing adaptive J2ME Applications Using AspectJ. In *Proceedings of the 7th Brazilian Symposium on Programming Languages*, pages 226–242, May 2003.
- [3] Brian Foote and Joseph Yoder. Metadata and Active Object-Models. Collected papers from the PLoP '98 and EuroPloP '98 Conference Technical Report wucs-98-25, Dept. of Computer Science, Washington University, September 1998.
- [4] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1994.
- [5] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification*. Addison-Wesley, second edition, 2000.
- [6] Jan Hannemann and Gregor Kiczales. Design Pattern Implementation in Java and AspectJ. In *Proceedings of the 17th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 161–173. ACM Press, 2002.
- [7] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William Griswold. Getting Started with AspectJ. *Communications of the ACM*, 44(10):59–65, 2001.

- [8] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Videira Lopes, Jean-Marc Loingtier, and John Irwin. Aspect-Oriented Programming. In *European Conference on Object-Oriented Programming, ECOOP'97*, LNCS 1241, pages 220–242, Finland, June 1997. Springer-Verlag.
- [9] Kalle Lyytinen and Youngjin Yoo. Issues and Challenges in Ubiquitous Computing: Introduction. *Communications of the ACM*, 45(12):62–65, 2002.
- [10] Gail C. Murphy, Robert J. Walker, Elisa L. A. Baniassad, Martin P. Robillard, Albert Lai, and Mik A. Kersten Kersten. Does Aspect-Oriented Programming Work? *Communications of the ACM*, 44(10):75–77, 2001.
- [11] Peyman Oreizy, Michael M. Gorlick, Richard N. Taylor, Dennis Heimbigner, Gregory Johnson, Nenad Medvidovic, Alex Quilici, David S. Rosenblum, and Alexander L. Wolf. An Architecture-Based Approach to Self-Adaptive Software. *IEEE Intelligent Systems*, 14(3):54–62, May 1999.
- [12] Vartan Piroumian. *Wireless J2ME Platform Programming*. Sun Microsystems Press, 2002.
- [13] Sérgio Soares and Paulo Borba. PaDA: A Pattern for Distribution Aspects. In *Second Latin American Conference on Pattern Languages of Programming, SugarLoafPLoP'2002*, pages 87–100, Itaipava, Brazil, 5th–7th August 2002.
- [14] Robert J. Walker, Elisa L. A. Baniassad, and Gail C. Murphy. An Initial Assessment of Aspect-Oriented Programming. In *Proceedings of the 21st International Conference on Software Engineering*, pages 120–130. IEEE Computer Society Press, 1999.
- [15] Joseph W. Yoder, Federico Balaguer, and Ralph Johnson. Architecture and Design of Adaptive Object-Models. *ACM SIGPLAN Notices*, 36(12):50–60, 2001.
- [16] Joseph W. Yoder and Ralph Johnson. The Adaptive Object-Model Architectural Style. In *Working IEEE/IFIP Conference on Software Architecture 2002(WICSA)*, Montréal, Québec, Canada, August 25-31 2002.