

Aspect-oriented approach to early design modelling

R. France, I. Ray, G. Georg and S. Ghosh

Abstract: Developers of modern software systems are often required to build software that addresses security, fault-tolerance and other dependability concerns. A decision to address a dependability concern in a particular manner can make it difficult or impossible to address other concerns in software. Proper attention to balancing key dependability and other concerns in the early phases of development can help developers better manage product risks through early identification and resolution of conflicts and undesirable emergent behaviours that arise as a result of interactions across behaviours that address different concerns. The authors describe an aspect-oriented modelling (AOM) approach that eases the task of exploring alternative ways of addressing concerns during software modelling. The paper focuses on use of the AOM approach to produce logical, aspect-oriented architecture models (AAMs) that describe how concerns are addressed in technology-independent terms. An AAM consists of a set of aspect models and a base architecture model called the primary model. An aspect model describes how a dependability concern is addressed, and a primary model describes how other concerns are addressed. Composition of the aspect and primary models in an AAM produces an integrated view of the logical architecture described by the AAM. Composition can reveal conflicts and undesirable emergent properties. Resolving these problems can involve developing and analysing alternative ways of addressing concerns. Localising the parts of an architecture that address pervasive and nonorthogonal dependability concerns in aspect models allows developers to more easily evolve and replace the parts as they explore alternative ways of balancing concerns in the early stages of development.

1 Introduction

The pervasiveness of computer systems highlights the need to engineer software that delivers services in a dependable manner. Designs of dependable software must address multiple, possibly interdependent, dependability concerns such as access control, confidentiality and data integrity. The manner in which a dependability concern is addressed can affect how other concerns are addressed. Balancing concerns during software development can involve developing and analysing alternative ways of addressing the concerns. Lack of attention to balancing dependability and other concerns in the early software development phases can lead to major rearchitecting of the design in later stages of development.

In this paper, a concern is a problem coupled with a desired goal [1, 2], where the goal determines acceptable solutions to the problem. For example, the problem of prohibiting unauthorised access to protected resources in a banking system is a dependability concern that must be addressed by banking software that manipulates the protected resources. A model that describes how a concern is addressed is called a *concern solution model*. In particular, a model that describes how a dependability concern is addressed is called a *dependability solution model*. For example, a role based access control (RBAC)

model [3] can be used to describe a solution to the banking system's access control concern. A decision to address a concern in a particular manner can give rise to other concerns. For example, the RBAC solution to the access control problem gives rise to new concerns pertaining to the management of roles and permissions.

This paper focuses on addressing dependability concerns during logical architecture modelling of software. The concern solution models are expressed in high-level, technology-independent terms. Current software development techniques allow developers to structure logical architectures in terms of modules that can be composite classes (i.e. classes that have an internal class structure), subsystems or interfaces. These modules typically localise solutions that address key functional concerns. Addressing nonorthogonal dependability concerns results in dependability solutions that are spread across the modules of the architecture and tangled with functionality described in the modules. These solutions are said to *crosscut* the primary structure of the architecture model.

Balancing concerns that are addressed by crosscutting solutions in the early phases of development can be challenging, primarily because of the difficulty of consistently changing or replacing the crosscutting solutions in an architecture model. A modelling approach that supports localising the descriptions of crosscutting dependability solutions can significantly ease the task of evolving and replacing the solution descriptions in an architecture model. In this paper we describe an aspect-oriented modelling (AOM) approach that allows developers to conceptualise, describe and communicate logical dependability solutions in isolation. The dependability solution models are called aspect models. An aspect-oriented architecture model (AAM) produced by the AOM approach consists of a set

of aspect models and a base architecture model called the primary model. The primary model describes concern solutions that determine the base structure of the architecture model. Each aspect model describes a dependability solution that crosscuts the primary model. An integrated view of the architecture is obtained by composing aspect and primary models to produce a composed AAM. Conflicts and undesirable emergent properties can be identified during composition of aspect and primary models and during analysis of the composed AAM. Addressing these deficiencies can lead to consideration of alternative ways of addressing concerns. Use of the AOM approach in the early stages of software development can help reduce software product risks through early identification and resolution of conflicts and undesirable behaviours that emerge as a result of integrating concern solutions.

2 Aspect-oriented modelling

In the aspect-oriented programming (AOP) language AspectJ, an aspect is a type that crosscuts a program structure [4]. An aspect contains information typically found in a class (i.e. data members and methods) in addition to behaviour that is executed at a specified point in a program's execution. The well defined points are called join points and the specifications of join points are called pointcuts. In the modelling community there has been some work on describing aspect-oriented programs using modelling languages such as the Unified Modelling Language (UML) [5]. AOM, as described in this paper, is not concerned with describing aspect-oriented programs. Rather, the AOM approach described in this paper provides support for modelling of concern solutions in isolation and for integrating the concern solution models with models describing the primary structure of software.

Modelling languages such as the UML provide some support for multidimensional separation of concerns through the use of different diagram types that can be used to describe nonorthogonal views of a system. AOM approaches allow developers to define additional dimensions of separation based on system-specific concerns. In an AOM approach, aspects localise concern solutions that crosscut views described by different diagrams in a system model.

The separation of crosscutting elements is a characteristic that is common to AOP and AOM, but differences between the artefacts (models versus code) can give rise to differences in techniques. For example, at the code level there is a single representation of functionality (the source code), while a model can describe a system from multiple views using different diagrams. The views can be nonorthogonal; for example, a UML sequence diagram that describes how a set of class instances interact to accomplish a task crosscuts the class diagram view of a system. In the AOM approach described in this paper, aspects describe solutions that crosscut UML model views.

Another difference between AOM and AOP is that code level aspect weaving is concerned primarily with inserting functionality at well defined points in a program's execution. The points at which functionality can be inserted are determined by the join point model of the AOP language. Software models are typically static descriptions of structure and behaviour. In the cases where the semantics of a modelling language supports execution of models, one can conceivably create a join point model for the modelling language to support an AOP-like notion of weaving. In the

absence of such semantics, weaving at the model level is essentially a static composition of model views.

2.1 Supporting aspect-oriented modelling

The AOM approach described in this paper provides support for (i) describing crosscutting concern solutions as modelling views called aspects, (ii) synthesising an integrated model by composing aspect and primary model views, and (iii) identifying and resolving conflicts and undesirable emergent properties that arise as a result of integrating aspect and primary models.

Two broad types of concerns can be identified [2]: A *concrete* concern has solutions that can be expressed in functional and structural terms in a model (i.e. there are model elements that specifically address the concern), and a *qualitative* concern is based on qualities or attributes of a system. Access control and error recovery are examples of concrete concerns, while concerns pertaining to system performance and memory utilisation are examples of qualitative concerns. The AOM approach described in this paper is applicable to concrete concerns only. Henceforth, a concrete concern is referred to simply as a concern.

Aspect models in our AOM approach describe crosscutting dependability solutions in logical (i.e. high-level and technology-independent) terms. A crosscutting concern solution can be isolated if its distributed elements have common structural and behavioural characteristics. A generalised form of the solution can then be represented as a pattern, where the pattern describes common characteristics of the distributed solution parts. A pattern view of crosscutting solutions screens out context-specific details and makes it possible to conceive, describe and understand the solutions in isolation. In our AOM approach an aspect model is a pattern that characterises a family of logical concern solutions. The patterns are described using UML model templates, as is also done in the Theme approach [6]. The template notation used in our work is an adaptation of a UML-based pattern language, called the Role-Based Metamodeling Language (RBML) [7]. Composing an aspect model with a primary model requires that one first instantiates the pattern by binding template parameters to application-specific values. An instantiated aspect model is called a *context-specific* aspect model. This approach paves the way for the development and systematic use of design patterns that capture logical solutions to dependability concerns.

Model composition technologies that automate significant parts of the AOM composition activity are needed if AOM is to scale-up to models of complex 'real-world' software systems. At one extreme are composition tools that take in aspect and primary models and produce composed models without further input from developers. This fixed composition approach provides very little flexibility in how aspect models are composed with primary models. At the other extreme, developers also provide composition procedures that detail how the aspect models are to be composed with primary models. This approach is very flexible, but requires more effort from developers. More practical solutions are likely to lie between these two approaches. For example, a tool can codify a default composition procedure and allow developers to vary some aspects of the procedure using composition directives. This is the approach taken in our work.

Context-specific aspect, primary and composed models are analysed to uncover flaws. Analysis of the composed model can reveal conflicts and undesirable emergent properties. Analysis can also be carried out to determine

the extent that dependability solutions meet their objectives when integrated with other concerns.

2.2 Overview of AOM approach

The major components of the AOM approach are shown in Fig. 1. An AAM of an application consists of (i) a primary model, (ii) aspect models and the bindings used to instantiate them in the application context, and (iii) composition directives that determine how the instantiated aspect models are composed with the primary model to produce a composed AAM.

A primary model consists of UML diagrams that each describe a view of the base architecture. The primary models in this paper consist of two types of diagram: UML classifier and interaction diagrams. Aspect models describe patterns of logical dependability solutions as UML diagram templates. An AAM presents logical views of a software architecture.

Figure 2 illustrates how an AAM consisting of two aspect models and a primary model is composed. The aspect models are instantiated by binding template parameters to application-specific values. We refer to the namespace from which binding values and names of elements in the primary

model are drawn as the *application domain namespace*. An aspect model can be instantiated multiple times to produce multiple context-specific aspects. Composition of context-specific aspect and primary models produces a model consisting of UML diagrams obtained by merging corresponding UML diagrams in the context-specific aspect and primary models. The AOM approach provides a basic composition procedure that can be altered in restricted ways by composition directives. For example, a composition directive can (i) specify that properties in aspect models override conflicting properties in primary models (or vice versa), (ii) specify that particular primary (or aspect) model elements must be removed or that new elements be added during composition, and (iii) determine the order in which two or more aspects are composed with a primary model.

The *Model Analysis* component in Fig.1 is responsible for analysing the composed model to identify errors and to determine the extent that dependability objectives are met. The focus of this paper is on aspect representation and model composition. We illustrate how identified conflicts can be resolved using composition directives, but a detailed account of techniques for analysing UML models is outside the scope of this paper.

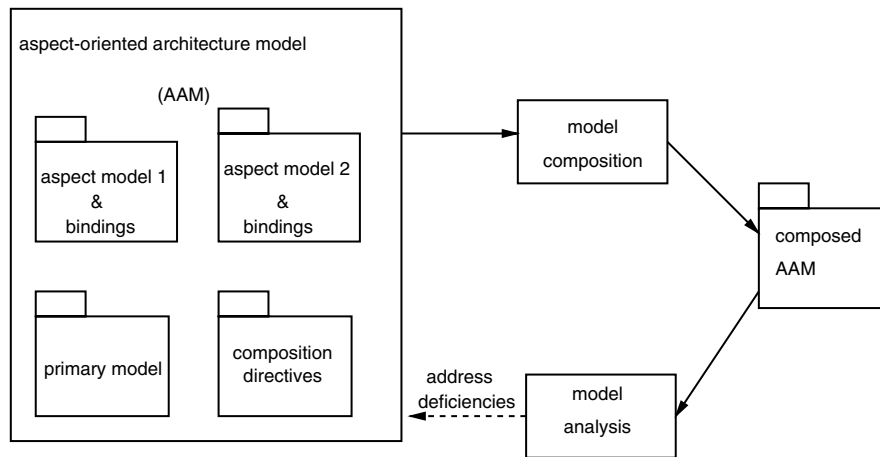


Fig. 1 Components of AOM approach

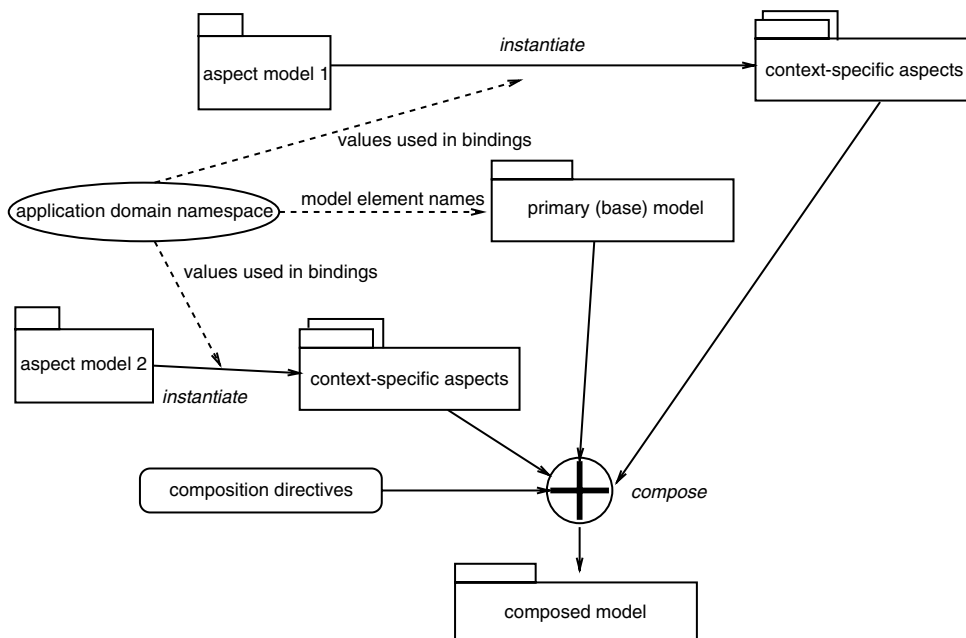


Fig. 2 Overview of composition in AOM approach

3 Representing aspect models

In this Section we describe how aspect models can be represented as template UML diagrams representing patterns of concern solutions. The template diagrams in this paper produce UML diagrams describing logical architectural views of solutions when instantiated.

In the UML, template models are described by parameterised packages that explicitly list the parameters in the package header. We have found this notation to be unwieldy when a large number of parameters are involved. In this paper the parameters are explicitly marked in the template diagrams using the symbol ‘|’.

Figure 3 shows an aspect model, *Auth*, characterising logical solutions in which access to a service is restricted to authorised clients. The aspect model consists of two diagram templates: a *class diagram template* that describes structural properties of the concern solutions, and a *collaboration diagram template* that describes interactions among solution elements. Instantiating the class diagram template shown in Fig. 3a results in a class diagram that consists of composite classes representing logical architectural views of clients, servers with services under access control, and authorisation repositories. A service under access control is represented by two operations in a server class:

- An operation that checks whether a client that requests the service is authorised to execute the service. The operation signature is obtained by instantiating the operation template *|operation*. The operation takes in as arguments the client’s identifier (represented by the operation argument template *|mid : |mgrid*) and zero or more values needed by the

service (represented by the argument template *|params1**). The template parameter *params1** is referred to as a *collection parameter*, indicating that it must be bound to a collection of values.

- An operation that performs the required service. This operation is obtained by instantiating the operation template *|doOperation*. The use of the *|params1** collection parameter in both the *operation* and *doOperation* templates indicates that the same value (i.e. the same set of arguments) must be used to instantiate the collection parameter in both of the templates.

The class template *|AuthorisationRepository* consists of the operation template *|checkAuth* that produces an operation that performs authorisation checks when instantiated. A *|checkAuth* operation uses the client identifier (represented by *|q : |mgrid*), an operation identifier (represented by *|op : |OpType*), and possibly other information passed in as arguments (represented by the collection parameter *|params2**), to determine whether the client is authorised to access the operation or not. If the client is authorised, then the operation returns a value that is an instantiation of *|valid*; otherwise it returns a value that is an instantiation of *|invalid*.

Operation templates may be associated with template forms of pre- and postconditions, referred to as *constraint templates*, that produce OCL specifications when instantiated. These constraint templates are presented separately from the diagrams to reduce diagram clutter. If an operation template is not associated with a constraint template, then an operation produced by the template must be specified in the primary model or its behaviour is to be specified or




Fig. 3 Authorisation-based access control aspect model

a Class diagram template for an authorisation aspect model
 b Collaboration diagram template for an authorisation aspect model

implemented in a subsequent refinement or detailing of the logical model. The operation templates `|doOperation` and `|checkAuth` do not have constraint templates associated with them. The following is the commented constraint template associated with the `|operation` template. The notation is based on the Object Constraint Language (OCL) version 2 [8]:

```
Context |Server::|operation(|mid:|mgrid,
  (|p: |T)*):
Pre:
  -- This operation can be invoked at any
  time.
  true
Post:
  /* The service is carried out if and only
  if the client is authorised to invoke
  the service.*/
  let authmessage : OclMessage =
    |AuthorisationRepository^|
      checkAuth(|mid,|opid,|p*) in
  (authmessage.hasReturned() and
    authmessage.result() = valid
    implies |Server^|doOperation(|p*)) and
  (|Server^|doOperation(|p*) implies
    authmessage.hasReturned() and
    authmessage.result() = valid)
```

The collaboration diagram template shown in Fig. 3b consists of template forms of participants (e.g. `:|Client`) and messages (e.g. `n:|operation(|id,|opParams*)`). An instantiated participant template produces either a named or anonymous participant; for example, binding `UserMgmt` to the parameter `Server` in the `:|Server` participant template produces the anonymous participant `:UserMgmt`. In a participant template, the type parameter (e.g. `|Server` in `:|Server`) must be a classifier template in a corresponding classifier diagram template. Participant type parameters and

the corresponding classifier templates must be instantiated with the same value.

Message templates consist of parameterised message sequence expressions and parameterised message expressions. For example, `|n.1: |auth := |checkAuth(|id, |opid, |other*)` consists of a parameterised sequence expression, `|n.1`, in which `n` is a parameter that can be substituted by a sequence expression (e.g. substituting 2.1.3 for `n` gives the sequence expression 2.1.3.1), and a parameterised message expression `|auth := |checkAuth(|id, |opid, |other*)` with parameters `auth`, `checkAuth`, `id`, `opid` and an optional set of arguments indicated by the collection parameter `other*`. The message expression `response := IDcheck(userid, updateOp, userstatus, usersession)` can be obtained from this template by binding `response` to `auth`, `IDcheck` to `checkAuth`, `userid` to `id`, `updateOp` to `opid` and `{userstatus, usersession}` to `other*`.

The collaboration diagram template for the `Auth` aspect model describes the following interaction pattern:

- Message `|n`: A client requests a service on a server by calling an instantiation of `|operation`.
- Message `|n.1`: An authorisation check is requested by calling an instantiation of `|checkAuth` in the authorisation repository linked with the server.
- Message `|n.2A`: If authorisation is granted, then the service is performed by invoking an instantiation of `|doOperation`.
- Message `|n.2B`: If authorisation is not granted the client is informed that access is not allowed.

The aspect model shown in Fig. 3 produces concern solution models that can be integrated with architecture models in which modules are composite classes (see [9] for more details on UML composite classes). Logical architectures can also be described using UML subsystems and interfaces as modules. The access control solution expressed in terms of subsystem templates is shown in Fig. 4a, and the logical




Fig. 4 Examples of subsystem and interface based access control aspect models

- a Subsystem diagram template for an authorisation aspect model
- b Interface diagram template for an authorisation aspect model

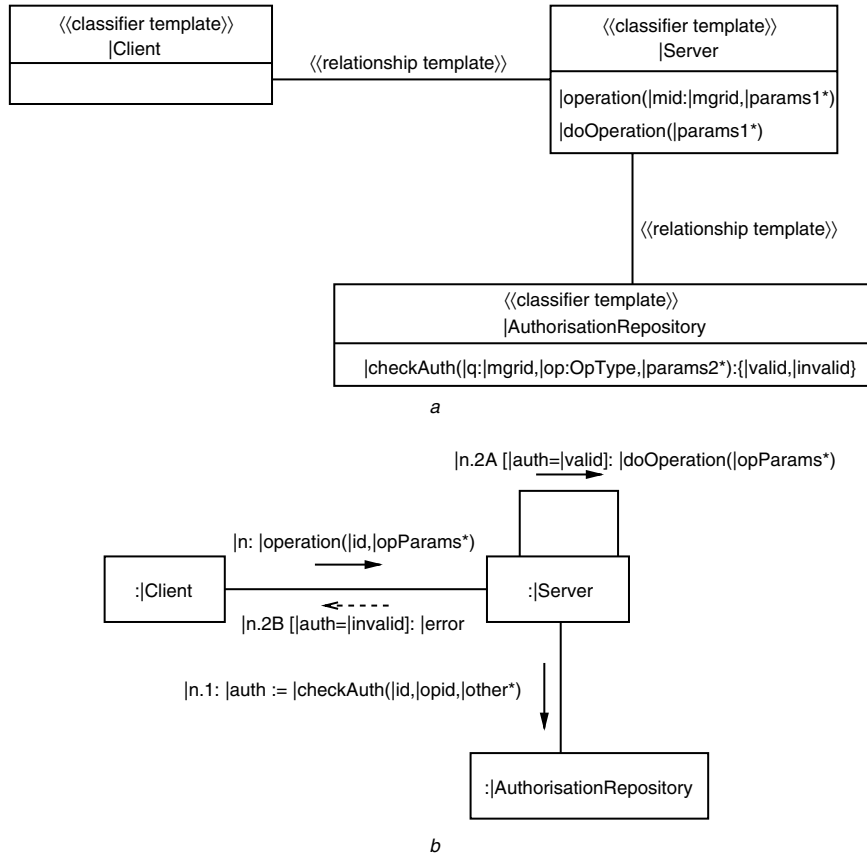


Fig. 5 Generalised access control aspect model

a Classifier diagram template for an authorisation aspect model
b Collaboration diagram template for an authorisation aspect model

solution expressed in terms of interface templates is shown in Fig. 4b. The collaboration diagram templates in these aspect models are syntactically identical to the collaboration diagram template shown in Fig. 3b and thus are not shown.

The three access control aspect models shown in Figs. 3 and 4 are specialisations of the aspect model shown in Fig. 5. The collaboration diagram of the generalised aspect is syntactically identical to the collaboration diagram shown in Fig. 3b. The generalised aspect model cannot be directly instantiated because it is based on abstract UML constructs (classifiers and relationships). This type of aspect model is called an *abstract* aspect model. An abstract aspect model must be specialised to a concrete aspect model (i.e. one based on concrete UML constructs) before it can be instantiated.

The remainder of this paper uses architecture models in which modules are composite classes to illustrate the AOM approach. The internal structures of the composite classes are hidden in the architectural views presented in this paper.

4 Composing aspect and primary models

Composing an aspect model with a primary model involves (i) instantiating the aspect model, using bindings, to produce a context-specific aspect model, and (ii) integrating the context-specific aspect model with the primary model. In this Section we illustrate how composition can be carried out using a small example.

4.1 Composition example

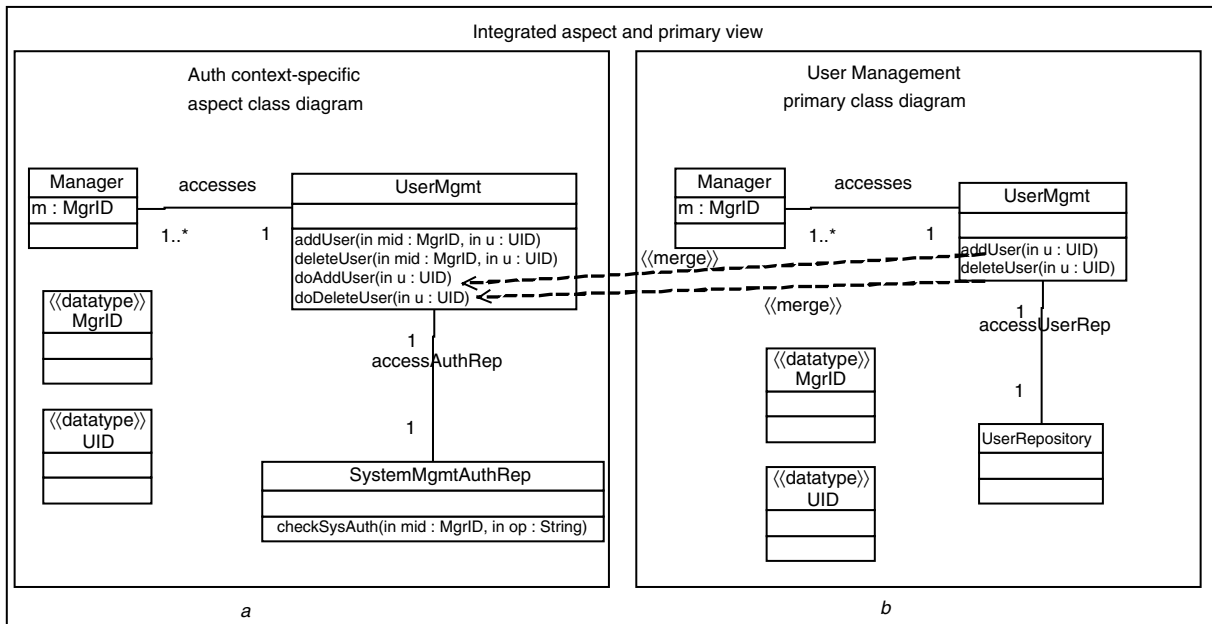
Figure 6b shows a primary model that describes a user management system in which *Manager* objects are linked to a *UserMgmt* object that controls access to a repository of

user information (a *UserRepository* object). The *UserMgmt* class defines operations for adding a user to the repository (*addUser*) and for deleting a user from the repository (*deleteUser*). Access to the *addUser* and *deleteUser* operations by *Manager* objects is unrestricted in the primary model. To restrict access to these operations the instantiated *Auth* aspect model shown in Fig. 6a is composed with the primary model to obtain the composed model shown in Fig. 6c.

The context-specific aspect model in Fig. 6a is obtained by instantiating the *Auth* aspect model using bindings that define the values that are to be substituted for parameters in the *Auth* diagram templates. A binding relates an aspect model element to a model element and can be expressed as a pair of the form (*aspect element name, model element name*). The *model element name* can be the name of a primary model element or the name of an application-specific element that is to be added to the composed model during composition. The type of the construct named by *model element name* must be the same as the parameter type; for example, a class template can only be bound to a model element that is a class. Some of the bindings used to produce the context-specific aspect model shown in Fig. 6a are given below:

```
(|Client, Manager); (|mgrid, MgrID);
(|accesses, accesses); (|m, 1..*); ((|n, |p, |q),1)[Note 1];
(|doOperation, doDeleteUser),
(|doOperation, doAddUser);
(|Server, UserMgmt); (|AuthorisationRepository,
SystemMgmtAuthRep).
```

Note 1: This is an abbreviated form of three pairs that respectively map n, p and q to the multiplicity 1.



Composition directives

Rename Primary::UserMgmt::addUser() to doAddUser()

Rename Primary::UserMgmt::deleteUser() to doDeleteUser()




Fig. 6 Example of composing a context-specific aspect class diagram and a primary class diagram

Note that a single parameter may be instantiated more than once; for example, the operation template *doOperation* is instantiated twice to produce *doAddUser* and *doDeleteUser* operations. An *instantiation multiplicity* can be associated with a template to restrict the number of times a template can be instantiated; for example, a template of the form *Temp 1..1* indicates that *Temp* can only be instantiated once. If a template is not associated with a (instantiation) multiplicity, then the number of instantiations possible is not restricted (as is the case in the aspect models given in this paper).

Sometimes it is more convenient to express bindings as relationships between structures. For example, the bindings for operation templates can be expressed as follows:

```
((operation, |mid, |params1*|, (addUser, mid, {u:UID}));
((operation, |mid, |params1*|, (deleteUser, mid, {u:UID}));
((checkAuth, |q, |op, |OpType, |params2*|,
(checkSysAuth, mid, op, String, {})).
```

Bindings also determine how constraint templates are instantiated. For example, the above bindings are used to

produce the following OCL definitions of the *addUser* and *deleteUser* operations in the context-specific aspect model:

```
Context UserMgmt :
  addUser (mid : MgrID, u : UID) :
Pre:
  true
Post:
  /* doAddUser() is called if and only if
  the Manager object is authorised to add
  users.*/
  let authmessage : OclMessage =
    SystemMgmtAuthRep^
      checkSysAuth(mid, ?:String) in
    (authmessage.hasReturned() and
    authmessage.result() = True
    implies self^doAddUser(u)) and
```

```

(self^doAddUser(u) implies
  authmessage.hasReturned() and
  authmessage.result() = True)
Context UserMgmt::
  deleteUser(mid:MgrID,u:UID):
Pre:
  true
Post:
  /* doDeleteUser() is called if and only if
  the Manager object is authorised to
  delete users.*/
let authmessage : OclMessage =
  SystemMgmtAuthRep^
  checkSysAuth(mid,?:String) in
(authmessage.hasReturned() and
  authmessage.result() = True
  implies self^doDeleteUser(u)) and
(self^doDeleteUser(u) implies
  authmessage.hasReturned() and
  authmessage.result() = True)

```

The AOM approach uses a basic name-based composition procedure in which elements with the same name are merged to form a single diagram element in the composed model. For example, merging the aspect and primary class diagram views of the *Manager* class results in a class that integrates information from both views. Some of the rules that determine how information associated with matching elements is combined are given below (these rules can be modified using composition directives, as indicated below):

- If the matching elements are operations with operation specifications, the operation specification in the composed model is the conjunction of the operation specifications associated with the matching operations. A composition directive can be used to vary how the specifications are logically connected.
- If the matching elements are attributes (or other elements) with constraints, the constraint associated with the attribute in the composed model is the conjunction of the constraints associated with the matching attributes. A composition directive can be used to vary how the constraints are logically connected.
- If the matching elements are associations, then the stronger (more restrictive) multiplicity at an association end is used in the composed model. A composition directive can be used to override this rule.

Unmatched model elements (i.e. model elements that only occur in either the aspect model or the primary model) are included in the composed class diagram.

Using the basic composition procedure to compose the *UserMgmt* aspect and primary model views results in a conflict for the *addUser* and *deleteUser* operations because they have the same names but different specifications in the two views; the *addUser* and *deleteUser* in the context-specific aspect model carry out authorisation checks, while the operations with the same names in the primary model add and delete users, respectively (the primary model specifications are not given in this paper). Furthermore, the operations *doAddUser* and *doDeleteUser* in the context-specific aspect model have the same specifications as those provided for *addUser* and *deleteUser* operations, respectively, in the primary model. Composition directives are used to resolve the conflict by renaming the *addUser* and *deleteUser* operations in the primary model to *doAddUser*

and *doDeleteUser*. The renaming removes the conflict and allows the primary model operations to be merged with the *doAddUser* and *doDeleteUser* operations in the context-specific aspect model. The result is that the *doAddUser* and *doDeleteUser* operations in the context-specific aspect model are respectively merged with the original *addUser* and *deleteUser* operations in the primary model as shown in Fig. 6.


In general, a composition directive can (i) determine the order in which multiple aspect models are composed with a primary model, (ii) define precedence or override relationships between matching aspect and primary model elements with conflicting properties and (iii) determine the elements that are renamed (e.g. to resolve conflicts), added, or deleted during composition. Adding new elements or deleting existing elements may be necessary to correctly compose aspect and primary models. For example, a security access control aspect may restrict access to an object by prohibiting particular relationships between the object and other objects. This can be done by identifying undesirable relationships in the aspect models and deleting them if found in the primary model. Elements marked for deletion in an aspect model are referred to as *prohibited elements*. Later in this Section we give an example of a situation that requires composition directives that add and delete model elements.

In summary, composition directives allow one to vary how aspect and primary models are composed. Consequently, aspect models do not need to capture all possible variations. In Section 4.2 we show how composition directives can be used to obtain variants of solutions described by aspect models.

4.2 Using composition directives to obtain variants of composed models

Using the same aspect and primary models, different composed models can be produced by varying the bindings and composition directives. Figure 7 shows a composed model obtained by composing the *Auth* aspect class diagram and the User Management primary class diagram using a different set of bindings and composition directives. We do not give the bindings for this case; they can be inferred from the context-specific aspect model shown in Fig. 7a.

The two systems described by the composed AAMs shown in Fig. 6c and in Fig. 7c accomplish the same tasks but do so differently. In Fig. 7c the authorisation operations and the services are located in separate classes. Rather than treating the *UserMgmt* class as a *Server* class, a new server class is introduced by the context-specific aspect model (see Fig. 7a)). The intent is that the *addUser* and *deleteUser* operations in the *UserAuth* class would call the corresponding operations in *UserMgmt* after a successful authorisation. To create a class diagram that reflects this intent, composition directives are defined that (i) add an association between the *UserAuth* and the *UserMgmt* classes, (ii) remove the association between the client and the *UserMgmt* class, (iii) remove the *doAddUser* and *doDeleteUser* operations from the *UserAuth* class, and (iv) replace references to *doAddUser* and *doDeleteUser* with references to *addUser* and *deleteUser*, respectively, in *UserMgmt*. The first directive is depicted by the association between the classes in the aspect and the primary model shown in Fig. 7. The second, third and fourth directives are captured by the following expressions:



Composition directives

Replace Primary::Manager::accesses by Aspect::Manager::uaccesses

Replace Aspect::UserAuth::doAddUser() by Primary::UserMgmt::addUser()

Replace Aspect::UserAuth::doDeleteUser() by Primary::UserMgmt::deleteUser()

prohibited elements are marked with an X

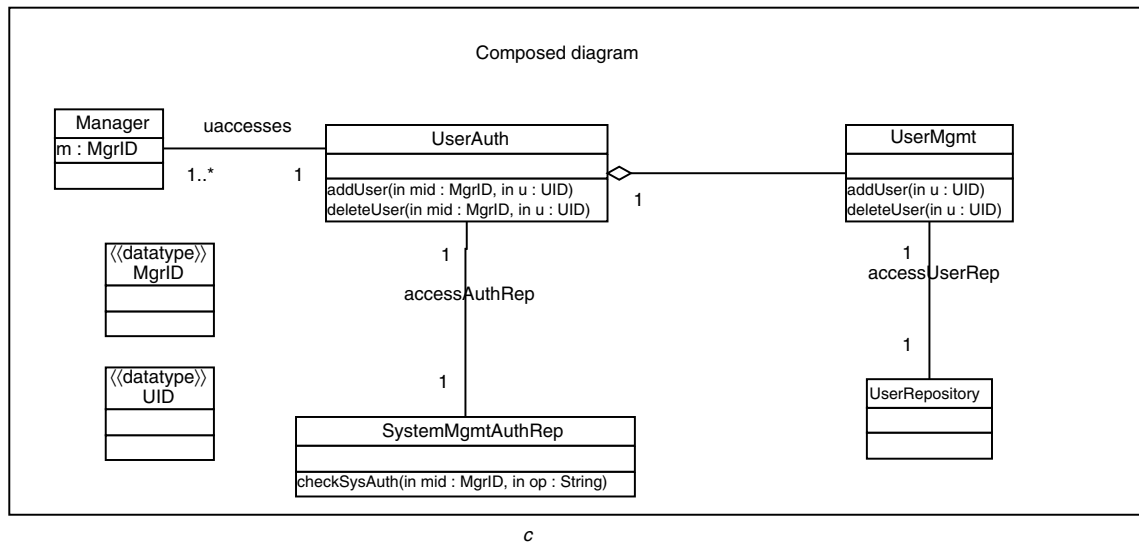


Fig. 7 Alternative composition of Auth and User Management class diagrams

Replace Primary::Manager::accesses by Aspect::Manager::uaccesses

- removes the accesses association between Manager and UserMgmt in the primary model (graphically indicated by placing an X on the association in the class diagram) and replaces all references to the association in Manager to the uaccesses association in the aspect model.

Replace Aspect::UserAuth::doAddUser() by Primary::UserMgmt::addUser()

- removes doAddUser (graphically indicated by placing an X on the operation in the class diagram) and replaces all references to it by references to addUser in the primary model.

Replace Aspect::UserAuth::doDeleteUser() by Primary::UserMgmt::deleteUser()

- removes doDeleteUser and replaces all references to it by references to addUser in the primary model.

The replacement of references is needed to ensure that the constraint definitions that refer to the deleted elements refer to their replacements in the composed model. For example, the above directives produce a composed model that includes the following operation definitions for *addUser* and *deleteUser* in *UserAuth* (this is obtained by replacing references to $\text{self}^{\wedge}\text{doAddUser}(u)$ by $\text{UserMgmt}^{\wedge}\text{addUser}(u)$, and $\text{self}^{\wedge}\text{doDeleteUser}(u)$ by $\text{UserMgmt}^{\wedge}\text{deleteUser}(u)$ in the OCL definitions of *addUser* and *deleteUser* given earlier for the primary model in Fig. 6b):

```

Context UserAuth::
  addUser (mid:MgrID,u:UID) :
Pre:
  true
Post:
  /*UserMgmt.addUser() is called if and
  only if the Manager object
  is authorised to add users.*/
  let authmessage : OclMessage =
    SystemMgmtAuthRep^
      checkSysAuth (mid,?:String) in
  (authmessage.hasReturned() and
  authmessage.result() = True
  implies UserMgmt^addUser(u)) and
  (UserMgmt^addUser(u) implies
  authmessage.hasReturned() and
  authmessage.result() = True)

Context UserAuth::
  deleteUser (mid:MgrID,u:UID) :
PreCondition:
  true
PostCondition:
  /*UserMgmt deleteUser() is called if and
  only if the Manager object is authorized
  to delete users.*/
  let authmessage : OclMessage =
    SystemMgmtAuthRep^
      checkSysAuth (mid,?:String) in
  (authmessage.hasReturned() and
  authmessage.result() = True
  implies UserMgmt^deleteUser(u)) and
  (UserMgmt^deleteUser(u) implies
  authmessage.hasReturned() and
  authmessage.result() = True)

```

In summary, we have shown how aspect and primary models can be composed and how composition directives can be used to resolve conflicts. One can view primary models and context-specific aspect models as views of an architecture, and thus their composition can be considered to be a view composition activity. We also show how composition directives and bindings can be used to produce different composed models from the same aspect and primary models. The example illustrates how bindings and composition directives can be used to reflect architectural decisions.

5 Limitations and open issues

In this Section we discuss some of the issues that are not yet addressed by our AOM approach, and outline our plans for addressing the issues.

5.1 Identifying aspects

It may not be desirable to model all crosscutting concern solutions as aspects. An AOM approach should provide guidelines that help developers determine the crosscutting concern solutions that can beneficially be localised in aspects. Our AOM approach targets crosscutting dependability solutions that may need to be balanced against other concern solutions, and those that are expected to evolve significantly during development. The localisation of these solutions can ease evolution of the solutions and provide support for rigorous tradeoff analysis. Currently, our AOM approach does not provide a set of detailed guidelines for

determining the crosscutting solutions that should be localised as aspects. Good guidelines should be based on experience and data collected on projects that utilise the AOM approach. Such experience and data are not yet available.

5.2 Developing composition strategies

When multiple aspect models are composed with a primary model, one has to be concerned with (i) the order in which the aspect models are composed and (ii) identifying and resolving conflicts or compromised behaviours. A conflict arises when a property in one aspect model contradicts a property in another aspect model. Composing a single aspect model with a primary model can also result in conflicts that need to be resolved (as shown in the preceding Section).

A behavior defined by an aspect model is compromised when it cannot be performed as specified because some of its sub-behaviours have been modified (or deleted) after merging with behaviours defined in other aspect models or the primary model. For example, (i) an aspect model may remove a relationship between two entities that is needed by a behaviour defined in another aspect model, or (ii) an operation replacement introduced by a composition directive results in behaviour that violates requirements previously satisfied by the operation being replaced. These problems can be resolved by making tradeoffs based on the relative importance of satisfying the conflicting requirements. Resolving problem (i) requires one to tradeoff the requirement that needs the relationship against the requirement that necessitates its deletion. Problem (ii) can be resolved by restoring the overridden operation and renaming the operation replacement.

It may be possible to apply prior experience in addressing concerns to constrain composition such that the occurrences of conflicts and compromised behaviours are minimised. Such experience can be captured in *composition strategies*. A composition strategy is influenced by domain knowledge pertaining to aspects (e.g. security and fault tolerance expertise), past experiences in addressing concerns, results of tradeoff analyses, and the properties (e.g. idempotency, commutativity, associativity and monotonicity) of the aspect models. Consider, for example, two security aspect models: one for authentication and the other for authorisation. Doing authorisation without authentication is meaningless. To get the desired result, an authentication aspect model must be composed with a primary model before an authorisation aspect model.

In summary, a composition strategy should be based on the properties of the aspect models, the constraints imposed by the domains of the aspect models, the results of tradeoff analyses, and the past experiences based on realising multiple, competing aspect models. A challenge is to develop a language for expressing composition strategies and techniques for obtaining composition directives from strategies. We are currently addressing these problems in our AOM research.

5.3 Analysing composed models

For large complex systems, the result of composition may be a complex model that may be difficult to comprehend. On the other hand, the composed model provides the detail needed to identify conflicts and undesirable emergent properties that arise as a result of interactions between model elements described by aspect and primary models. In the AOM approach, composition is carried out primarily to support analysis that uncovers conflicts and other defects

that arise as a result of integrating aspect and primary model views.

Analysis can be performed at three levels: unit analysis is concerned with analysing a single context-specific aspect model, integration analysis occurs when context-specific aspect models are composed sequentially with a primary model, and system analysis occurs when all aspect models have been composed with a primary model. When multiple aspect models are composed sequentially with a primary model, one must test that no existing capabilities were broken and that the capability of the newly composed aspect was preserved in the composition. System analysis is concerned with determining whether the composed AAM satisfies the requirements.

To support dynamic evaluation of composed models, an operational semantics for UML models is needed. We are currently adapting a systematic technique for testing and exercising UML designs to our AOM approach [10–12]. State exploration techniques, such as model-checking (e.g. see [13, 14]), can also be used to analyse composed models.

We are also developing a system analysis technique that involves evaluating composed models against a representative set of usage scenarios. The scenarios describe both proper and improper uses of the system. Scenarios describing improper usages are called *misuse* scenarios. For example, to evaluate the impact of security concern solutions on an application, misuse scenarios that describe malicious attacks can be developed. Misuse scenarios are used to determine if the mechanisms defined by the security aspect models are sufficient to prevent the attacks from compromising protected resources. Scenarios describing authorised interactions are used to determine if the authorised activities are adversely affected by behaviours described by the aspect models. Scenarios are expressed in terms of UML behavioural models (e.g. sequence diagrams) and can be based on use cases that describe authorised behaviours and on misuse cases that describe behaviours that should not be present in a correct system implementation. Analysis involves composing the scenario descriptions with the composed AAM and evaluating the result. If correct composition of a misuse scenario and an AAM produces a consistent model, then the AAM has a flaw. Similarly, if correct composition of a scenario describing authorised interactions and an AAM produces an inconsistent model, then the AAM has a flaw.

5.4 Evolving aspect-oriented models

Support for extracting aspect and primary model views from composed models can help ease the task of evolving AAMs. For example, changing a context-specific aspect model or primary model after composition has been carried out can be accomplished by extracting the model view from the composed model and changing the model. Reintegration of the view involves propagating the changes to the other parts of the composed model. Similarly, developers should be able to add new composition directives or modify composition directives and bindings in order to resolve conflicts and fix other defects in the composed model. Extraction of model views or composition related information, and their reintegration, requires tool support if this approach is to scale-up to large system models. Automated support for view extraction and reintegration also eases exploration of solution alternatives carried out to support tradeoff analysis.

Extracting views and other information used to compose models requires maintaining relationships among aspect

models, the primary model and the composed model. The problem is similar to tracking the evolution of complex composite parts in discrete manufacturing. We are currently investigating the use of a standard framework, the product data management (PDM) framework [15], that was developed for managing the evolution of complex products in the discrete manufacturing area to support storing and evolving AAMs.

5.5 Process support for architectural modelling using AOM

AOM can be carried out in the context of an iterative and incremental architecture development process. In the first iteration an initial primary model is developed. This model reflects early decisions pertaining to the concerns that determine the modular structure of the architecture. The AOM approach described in this paper supports the development of an architecture model in which the modules can be composite classes, subsystems or logical components.

An initial set of context-specific aspects that describe logical solutions to a subset of dependability concerns that crosscut the primary model are also developed in the first iteration. The initial aspect and primary models are composed to produce a composed AAM, which is then analysed. The following activities are carried out in each subsequent iteration:

- If the analysis performed in the previous iteration uncovers problems, the aspect models, primary model or the composition directives are modified accordingly.
- New aspect models for dependability concerns not covered in previous iterations can be introduced, and the primary model can also be extended to take into account functionality not considered in previous iterations.
- If needed, new composition directives are created.
- The modified aspect and primary models are composed and analysed.

5.6 Tool support for AOM

We are developing a prototype integrated toolset that supports (i) creation and cataloging of aspect models, (ii) composition of aspects and primary models, and (iii) rigorous analysis of composed models. To date, our work on tool development has produced the following:

- An architectural design of a toolset that supports creation of aspect models, composition of aspects and primary models, and analysis of composed models has been developed [16].
- A prototype editor for creating aspect model class diagram templates has been developed. The editor was built using the Eclipse modelling framework (see <http://www.eclipse.org/emf>). The prototype does not support instantiation of the templates.
- A tool, built on top of Rational Rose, that generates instantiations from template forms of UML class diagrams (generic aspects) has been developed.
- A prototype model composer that takes primary model and context-specific aspect class diagrams and composes them has been developed.

We are currently integrating and extending the above tools to form an integrated AOM tool set.

6 Related research

Aspect-oriented programming (AOP) supports multi-dimensional separation of concerns (MDSoc) at the

programming level [4, 17–24] An AOP aspect is an implementation or design concern that crosscuts the primary functional units of a program (e.g. concerns that crosscut classes of an object-oriented program). A few researchers have started to address the problem of defining and composing aspects at an abstraction level higher than the programming language level (for example, see [6, 25–29].

Fiadeiro and Lopes [25] specify aspects related to system co-ordination using an algebraic approach. Their approach is applicable to detailed design and code, and utilises a notation that is not widely known by system developers. Gray *et al.* [26] use aspects to represent aspects in domain-specific models. Their research is part of the model-integrated computing (MIC) initiative that targets embedded software systems specifically. MIC extends the scope and usage of models such that they form the backbone of a development process for building embedded software systems. Requirements, architecture and the environment of a system is captured in the form of formal high-level models that allow the representation of concerns. Our work on MDSoC can complement the MIC efforts by providing UML-based techniques for representing and composing aspects, and making tradeoff decisions. Suzuki and Yamamoto [29] extend the UML so that it can be used to model code level aspects. Their approach is restricted to design aspects that can be represented as aspects in an aspect-oriented program.

In the AOM approach proposed by Clarke [30], and Clarke and Walker [31, 32], a design called a *Subject* is created for each system requirement. A comprehensive design is a composition of subjects. Subjects are expressed as UML model views. Composition relationships specify how models are to be composed by identifying overlapping concepts in the subjects and specifying how models are integrated. The UML metamodel is extended to support composition relationships and describe well-formedness rules for composition. Two types of integration strategies are used: override and merge. Override integration is used when existing behaviour in a subject needs to be updated to reflect new requirements. Merge integration is used when subjects for different requirements are to be integrated. Operations in related subjects may need to be merged into a unified operation. Reconciliation strategies are used to resolve conflicts between property values of corresponding subject elements. Precedence relationships, transformation functions applied to conflicting elements, explicit specification of reconciled elements, and default values may be used for reconciliation.

As part of the Early Aspects initiative, Rashid and co-workers have targeted multi-dimensional separation throughout the software cycle [27, 28, 33, 34]. This work supports modularisation of broadly scoped properties at the requirements level to establish early tradeoffs, provide decision support and promote traceability to artefacts at later development stages.

The work described in this paper extends our previous work (for example, see [35–37]) by refining the aspect modelling notation and the instantiation process, and refining the notion of composition directives to support conflict resolution and modelling of solution variants.

7 Conclusion

Current modelling approaches provide good support for modularising systems along a few dimensions. AOM can significantly enhance support for separation of concerns targeted at tackling growing software complexity. The AOM approach described in this paper can help developers better manage the complexity of creating and evolving

complex software that must address multiple dependability concerns.

Our research goal is to develop an AOM approach that addresses three factors that contribute to the complexity of software development: (i) the complexity inherent in the required functionality of the software system; (ii) the pervasiveness and variety of interdependent concerns that must be addressed in an architecture; and (iii) the need to balance forces when addressing competing system concerns. The above factors can be addressed in an AOM approach that integrates work on model-driven development, MDSoC, and value-based assessment. Model-driven development addresses factor (i) by raising the level of abstraction at which functionality is developed. Approaches that support MDSoC address factor (ii) by providing the means for isolating, composing and analysing crosscutting solutions. Value-based assessment techniques address factor (iii) by providing a base for rigorous tradeoff analysis. The work described in this paper addresses factors (i) and (ii). We are currently developing support that explicitly addresses factor (iii). Tradeoff analysis is desired when crosscutting solutions interact in ways that compromise the accomplishment of concern objectives. In such situations the system developer must make tradeoffs based on prioritisations of objectives. The challenge is to (i) develop systematic and quantitative tradeoff analysis techniques that allow developers to assess alternative solutions, (ii) develop techniques for capturing and representing experience related to making tradeoffs across a set of aspect models, and (iii) use the captured experience to guide how aspect models are composed with other models. The captured experience can take the form of composition strategies that determine the set of aspect models and composition directives that produce a composed model that best meets the requirements. Composition strategies and the decisions they drive should be based on information about the value and importance of the architectural choices represented in alternative aspect models. Our ongoing work in this area involves adapting existing approaches to tradeoff analysis – for example, the DDP approach [38–40].

8 Acknowledgments

This material is based on work funded by AFOSR under award no. FA9550-04-1-0102.

9 References

- 1 Jackson, M.: 'Problem frames: analyzing and structuring software development problems' (Addison Wesley Professional, 2000)
- 2 Kande, M.: 'A concern-oriented approach to software architecture'. PhD thesis, EPFL, Lausanne, Switzerland, 2003
- 3 Sandhu, R., Coyne, E.J., Feinstein, H.L., and Youman, C.E.: 'Role based access control models', *Computer*, 1996, **29**, (2), pp. 38–47
- 4 Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., and Griswold, W.G.: 'An overview of AspectJ'. Proc. European Conf. on Object-Oriented Programming (ECOOP), Budapest, Hungary, June 2001, pp. 327–353
- 5 The Object Management Group. 'Unified modelling language. Version 1.5', OMG, formal/2003-03-01, 2003
- 6 Clarke, S., and Murphy, J.: 'Developing a tool to support the application of aspect-oriented programming principles to the design phase'. Proc. Int. Conf. on Software Engineering (ICSE), Kyoto, Japan, April 1998
- 7 France, R., Kim, D.-K., Ghosh, S., and Song, E.: 'A UML-based pattern specification technique', *IEEE Trans. Sofw. Eng.*, 2004, **30**, (3)
- 8 Warmer, J., and Kleppe, A.: 'The object constraint language' (Addison-Wesley, 2003, 2nd edn.)
- 9 The Object Management Group. 'Unified modelling language: superstructure. Version 2.0', OMG, ptc/03-07-06, 2003
- 10 Andrews, A., France, R., Ghosh, S., and Craig, G.: 'Test adequacy criteria for UML design models', *Sofw. Test. Verif. Reliab.*, 2003, **13**, (2), pp. 95–127
- 11 Ghosh, S., France, R.B., Braganza, C., Kawane, N., Andrews, A., and Pilskalns, O.: 'Test adequacy assessment for UML design model

- testing'. Proc. Int. Symp. on Software Reliability Engineering (ISSRE), 2003
- 12 Pilskalns, O., Andrews, A., Ghosh, S., and France, R.B.: 'Rigorous testing by merging structural and behavioral UML representations'. Proc. 6th Int. Conf. on Unified Modeling Language (UML), San Francisco, USA, 20–24 October 2003
 - 13 Atlee, J.M., and Gannon, J.: 'State-based model checking of event-driven system requirement', *IEEE Trans. Softw. Eng.*, 1993, **19**, (1), pp. 24–40
 - 14 Holzmann, G.J.: 'Design and validation of computer protocols' (Prentice Hall, Englewood Cliffs, NJ, 1991), Software Series
 - 15 O.M. Group. 'PDM enablers'. Technical Report mfg/98-02-02, The Object Management Group (OMG), 2002
 - 16 Mekerke, F., Georg, G., France, R., and Alexander, R.: 'Tool support for aspect-oriented design'. Presented at Workshop on Advances in Object-Oriented Information Systems: (OOIS), 2002
 - 17 Bergmans, L., and Aksit, M.: 'Composing multiple concern using composition filters', *Commun. ACM*, 2001, **44**, (10)
 - 18 Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., and Grisword, W.G.: 'Getting started with AspectJ', *Commun. ACM*, 2001, **44**, (10), pp. 59–65
 - 19 Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C.V., Loingtier, J.-M., and Irwin, J.: 'Aspect-oriented programming', *Lect. Notes Comp. Sci.*, 1997, **1241**, pp. 220–242
 - 20 Kieberherr, K., Orleans, D., and Ovlinger, J.: 'Aspect-oriented programming with adaptive methods', *Commun. ACM*, 2001, **44**, (10), pp. 39–41
 - 21 Ossher, H., and Tarr, P.: 'Using multidimensional separation of concerns to (re)shape evolving software', *Commun. ACM*, 2001, **44**, (10), pp. 43–50
 - 22 Pace, J.A.D., and Campo, M.R.: 'Analyzing the role of aspects in software design', *Commun. ACM*, 2001, **44**, (10), pp. 66–73
 - 23 Silva, A.R.: 'Separation and composition of overlapping and interacting concerns'. Proc. 1st Workshop on Multi-Dimensional Separation of Concerns in Object-Oriented Systems (OOPSLA), Denver, Colorado, November 1999
 - 24 Sullivan, G.T.: 'Aspect-oriented programming using reflection and metaobject protocols', *Commun. ACM*, 2001, **44**, (10), pp. 95–97
 - 25 Fiadeiro, J.L., and Lopes, A.: 'Algebraic semantics of co-ordination or what is it in a signature?', *Lect. Notes Comp. Sci.*, 1999, **1548**, pp. 293–307
 - 26 Gray, J., Bapty, T., Neema, S., and Tuck, J.: 'Handling crosscutting constraints in domain-specific modeling', *Commun. ACM*, 2002, **44**, (10), pp. 87–93
 - 27 Rashid, A., Moreira, A., and Araujo, J.: 'Modularization and composition of aspectual requirements'. Proc. 2nd ACM Int. Conf. on Aspect-Oriented Software Development, Boston, March 2003, pp. 11–20
 - 28 Rashid, A., Sawyer, P., Moreira, A., and Araujo, J.: 'Early aspects: a model for aspect-oriented requirements engineering'. Proc. IEEE Joint Int. Conf. on Requirements Engineering, Essen, Germany, 9–13 September 2002, pp. 199–202
 - 29 Suzuki, J., and Yamamoto, Y.: 'Extending UML with aspects: aspect support in the design phase'. Proc. 3rd ECOOP Aspect-Oriented Programming Workshop, Lisbon, Portugal, June 1999
 - 30 Clarke, S.: 'Extending standard UML with model composition semantics', *Sci. Comput. Program.*, 2002, **44**, (1), pp. 71–100
 - 31 Clarke, S., and Walker, R.J.: 'Composition patterns: An approach to designing reusable aspects'. Presented at 23rd Int. Conf. on Software Engineering (ICSE), Toronto, Canada, 2001
 - 32 Clarke, S., and Walker, R.J.: 'Towards a standard design language for AOSD'. Presented at 1st Int. Conf. on Aspect-Oriented Software Development, Enschede, The Netherlands, April 2002
 - 33 Rashid, A.: 'A hybrid approach to separation of concerns: the story of SADES', *Lect. Notes Comp. Sci.*, 2001, **2192**, pp. 231–249
 - 34 Rashid, A., and Chitichyan, R.: 'Persistence as an aspect'. Proc. 2nd ACM Int. Conf. on Aspect-Oriented Software Development, Boston, March 2003, pp. 120–129
 - 35 Georg, G., France, R., and Ray, I.: 'An aspect-based approach to modeling security concerns'. Proc. Workshop on Critical Systems Development with UML, Dresden, Germany, 2002
 - 36 Georg, G., France, R., and Ray, I.: 'Designing high integrity systems using aspects'. Proc. 5th IFIP TC-11 WG 11.5 Working Conf. on Integrity and Internal Control in Information Systems (IICIS), Bonn, Germany, November 2002
 - 37 Georg, G., Ray, I., and France, R.: 'Using aspects to design a secure system'. Proc. Int. Conf. on Engineering Complex Computing Systems (ICECCS), Greenbelt, MD, December 2002
 - 38 Feather, M.: 'A quantitative risk-based model for reasoning over critical system properties'. Proc. Int. Workshop on Requirements for High Assurance Systems, Essen, Germany, September 2002, pp. 11–18
 - 39 Feather, M., Conford, S., Dunphy, J., and Hicks, K.: 'A quantitative risk model for early lifecycle decision making'. Presented at Society for Design and Process Science Conf. on Integrated Design and Process Technology (IDPT), 2002
 - 40 Feather, M., and Cornford, S.: 'Quantitative risk-based requirements reasoning', *Require. Eng.*, to be published