

Evolution of Aspects for Legacy System Security Concerns

Robin C. Laney
Dept. of Computing,
The Open University, UK
+44 (0)1908 654342
r.c.laney@open.ac.uk

Janet van der Linden
Dept. of Computing,
The Open University, UK
+44 (0)1908 652985
j.vanderlinden@open.ac.uk

Pete Thomas
Dept. of Computing,
The Open University, UK
+44 (0)1908 652695
p.g.thomas@open.ac.uk

ABSTRACT

This paper shows how aspects can be successfully employed in the support of system evolution. The context is a case study on migrating a legacy client-server application to overcome the security problems associated with ‘message tampering’ attacks. The focus is on authorization issues in which aspects are used to add a security mechanism based on digital signatures. The approach provides for future evolution of the system. In particular, it is shown how factoring of aspectual concerns allows the scope of the security boundary to be varied, illustrating reuse of the aspects.

Whilst the aspects are added non-intrusively, it is demonstrated how aspects can modify the control-flow behaviour of a server. We propose a programming idiom that conforms to Design by Contract.

Categories and Subject Descriptors

D.3.3 [Programming Languages]: Language Constructs and Features

General Terms

Design, Experimentation, Security, Standardization, Languages

Keywords

Aspects, security, legacy system

1 INTRODUCTION

The work reported on here arose from the need to extend a case study that our third year undergraduate students use in their study of distributed systems. The package is based around a home banking system written in Java that supports multiple distributed bank servers as well as many home users. It consists of a number of separate networked modules dealing with a variety of middleware issues including load balancing and routing. Whilst the system is not large by comparison with real home banking systems, it is large enough to present real problems in terms of scale and complexity. Over the years, the system has been extended and modified as new features were introduced and software engineering principles such as refactoring were employed. Our most recent requirement was to extend the security facilities within the system and we felt that it would be useful to

students if we were to approach the problem using aspects. Indeed, the package has been so extensively modified that it exhibits all the properties of legacy software, in particular: modifications to the original specification by different people with variable amounts of documentation of dubious quality. We felt that the time to make even small modifications would be prohibitive and we did not wish to change the existing course materials that describe the system. Therefore, extending the system in a non-intrusive manner should be a good test of the use of aspects and should enable us to continue to grow the system in a comprehensible way. Our aim is not to address all security issues, for example we do not address issues of Decentralized Trust Management [3]. In this paper we report on the problems we encountered, how the problems were overcome and the insights we gained into AOP.

A number of researchers [5, 15, 16] have identified reasons why aspects are a good mechanism for improving security in software systems:

- Current approaches to the development of security architectures for applications result in software that has application code tangled with security code [16]. Techniques that help to control and reduce this pervasiveness should be exploited.
- Security experts will be able to better focus their efforts on security issues if these are concentrated and separated from the rest of the application.

In [5] it is discussed how security requirements can be viewed through a number of questions: *how*, *what*, *where* and *when* should security issues be addressed in the application.. In the past, there have been successful attempts to factor out some of these questions. That is, *what* and *how* can be addressed, for example, through the use of libraries, for cryptography, or through a service such as JAAS, the Java Authentication and Authorization Service. However, the question of *where*, for example, should each security concern be implemented at the beginning of each method, or only in some methods? and the question of *when*, for example, should a concern be addressed when some application condition is satisfied or not? are both still tangled with the application. This is because the calls to the security libraries or services have to be made from within the application. Currently, aspects are beginning to address some of these *where* and *when* issues. For example, [11] discusses how to make calls to JAAS using aspects.

In our work we have used aspects to implement authorization requirements based on digital signatures. We were able not only to separate the lower level mechanism of implementing a digital signature, but also to factor out the policy of where we want this security concern to be applied.

The use of aspects in software development has been divided into two categories [2, 10]: development and production aspects. Development aspects can be used during the development of applications to facilitate debugging, testing and performance tuning, and are not part of the final build of the application. In our work we made use of logging aspects to assist in debugging. Production aspects are intended to be included in the build of an application, and provide additional functionality for the application. There is evidence that production aspects are helpful in factoring the design of green-field developments [1, 9], as well as for existing systems [3, 6]. In this paper we explore their application in the evolution of a legacy system, under the strict constraint that no changes are to be made intrusively.

Our work is based on the use of a case study encompassing multiple servers in a prototype home banking application. The existing system is secure to the degree that users must log-on before accessing an account. However, messages between clients and servers are not signed, exposing the system to attacks caused by message tampering. Our initial goal in this work was to add digital signatures to the system to enhance the authorization mechanisms. We focused on authorization issues simply to clarify the separation of concern issues. Having achieved this and solved some interesting but unexpected problems concerning the degree of interaction between an aspect and the legacy code, we turned our attention to future system evolution. We give a specific illustration of evolving the system, in the light of revised security requirements involving a transfer of responsibility for signing and checking messages between servers. We discuss the benefits of well designed aspects that are reusable in system evolution.

The paper is organized as follows. In section 2, we discuss related work and present relevant details of the case study. In section 3, we specify our initial additional security requirements and discuss the design and implementation necessary to meet them. In particular, we describe the difficulties of developing aspects that must modify control flow behaviour and how the exception mechanism is of assistance in solving this problem. Furthermore we show how our aspects were factored to enhance separation of concerns. In section 4, we introduce a requirement to move the security boundary and show how the factoring of our aspectual security code smoothes the evolution of the system. Section 5 draws conclusions and outlines future work.

2 BACKGROUND AND RELATED WORK

2.1 Related Work

Security is increasingly becoming an area that attracts the attention of the AOSD community. In many ways security represents an ideal candidate for the separation of concerns, because it permeates all areas of software systems. Security would thus benefit from techniques that help to structure and modularize it.

One example of using security aspects is to modularize access control. In [6] authentication is separated from the application, here an FTP server. Aspects were used to perform the task of checking username/password combinations, while other aspects represented the authenticated user thus holding important security information for the duration of an FTP session.

A different approach to modularizing access control is presented in [11]. Laddad shows how JAAS, the Java Authentication and Authorization Service can be added to applications using aspects, rather than making calls to this service from the application code.

Security issues can also relate to the notion of secure code: many programming languages have features that can easily be used in such a way as to leave security flaws. In [15] Viega, Bloch and Chandra discuss how aspects can help developers to be consistent in their implementation of security policies by, for example, checking that calls to the SecurityManager are included where needed, replacing all calls to non-secure functions by calls to secure functions, providing buffer overflow protection or logging data that may be relevant to security.

An alternative approach to implementing non-functional requirements is the use of reflection language systems/architectures. In particular Welch and Stroud [16] present a portable approach based on their Kava meta-object protocol. They present a case study in which a security system based on proxies and inheritance is re-implemented using Kava. This results in better enforcement of security but the use of reflection does raise issues of efficiency. Their approach is portable as it is not dependent on specialized language systems or architectures as is the case with most other work on reflection.

Other work has looked at the factorization of middleware software [17], based on a mining approach to discover scattered code in legacy systems that can be refactored as aspects. This is in contrast to the work presented in this paper where we are adding functionality.

2.2 The Case Study

The case study that we are working on was developed as a tool for illustrating concepts in concurrent programming and for illustrating software engineering techniques. As such, the software has been continually amended as new requirements were specified.

The system was based on the following scenario. A bank wished to provide its existing customers with the ability to perform a number of different banking transactions using their home computers. The bank started this development from the point at which a customer already had one or more accounts at one or more of its branches and each branch had a computer-based banking system to maintain its accounts. In addition, the branch computers communicated with a central machine at the bank's headquarters. The architecture of the home banking system is shown in Figure 1.




Figure 1 The home banking system

A Branch Server is an existing branch system: there are many branches. The Client represents a customer's machine. Between a Client and its Branch Server are three items of middleware. The Queue Server interfaces with all Clients and queues customer requests before passing them on to the Proxy Server. One of its purposes is to smooth out the flow of requests to the Proxy Server. The Proxy Server acts as a router, using the customer's account number to route the request to the appropriate branch. The Proxy Server uses the facilities of the Name Server to determine the IP address of the appropriate Branch Server. The home banking system was designed so that each of the servers could be hosted on separate machines. Indeed, there is provision for installing multiple Proxy Servers should the loading on a single server become too high. It was envisaged that the Queue, Proxy and Name Servers would be housed at the bank's headquarters. The shaded area illustrates that the Queue Server acts as an interface between the Clients and the Servers and deals with some security concerns.

Client requests are in the form of asynchronous messages packaged as strings constructed according to a proprietary protocol known to each server. Every message contains the customer's identifier and account number. Responses to client requests are also messages conforming to a specific protocol and are routed back from the Branch Servers via the Proxy and Queue Servers to the Client. The Queue Server is responsible for forwarding a response to the appropriate client.

In the initial prototype implementation, security is minimal: authentication is performed by the Branch Server using a single password. Client requests are rejected by the Queue Server unless the client has successfully logged-in. If the Queue Server receives a login request message it passes it directly to the appropriate Branch Server for authentication. The Branch Server returns a message to the QueueServer and the Client indicating whether or not the authentication succeeded. If the login is successful, the Queue Server will accept all future requests from that Client until the client ends the session.

3 Adding Digital Signatures

The principle of digital signatures is to calculate a signature based on a given message string, and send this signature along with the original message. On arrival the recipient can check whether the signature still matches the message. By adding a mechanism for signing messages using digital signatures we are able to reduce greatly the susceptibility of our system to 'message tampering' attacks. Our focus in this research is on authorization: the signing and checking of messages, ignoring the separate issue of authentication.

As was outlined above, messages are packaged as strings according to proprietary protocols. As messages are passed around from server to server, they are taken apart and built up again as they get processed. A small change in the protocol at one server, or client, will impact on all subsequent servers. At each point where messages arrive at a server, a thorough checking of the message format takes place, and all the action in the server depends on this precise format of the message. Altering a message format, by adding a digital signature would impact on the main code structure in each server. Therefore, aspects are relevant as an implementation technique, because we need to interact with the message passing code that crosscuts the system. An alternative approach based on the use of the Delegation or Proxy pattern would involve restructuring of the legacy code.

For reasons of traceability from requirements through design to implementation, it is important that we relate aspects structurally and functionally to the requirements. As we show below, these details emerge iteratively with design and implementation of the aspectual details. These issues, which we explore further below suggest that the use of early aspects i.e. at the requirements stage [7, 14], exploited within the context of a twin-peaks approach [13] aid the development of secure code.

3.1 Relating Requirements and Aspects

The requirement to improve authorization security through the use of digital signatures entails:

- (1) Generating and checking signatures at appropriate places, that is, at a security boundary, and storing user keys at sites where the signing of messages and the checking of signatures takes place.
- (2) Modifying message handling to allow incorporation of signatures in messages. In this system, the Queue Server, Proxy Server and Branch Server have legitimate reasons for viewing (some of) the contents of a message and therefore have to unmarshal and marshal the message.

The initial security boundary, that is the perimeter of an area encompassing those parts of the system that are to be secured, included all the servers and parts of the clients. The shaded area in Figure 2 illustrates this. Note that the aspects are not required to ensure security of the whole of the client. Prior to signing messages, security is necessarily the responsibility of some other combination of software/hardware.




Figure 2: Banking application with security boundary

Our initial design involved two aspects crosscutting the two requirements (in addition to crosscutting within, and as it emerged, across servers). The aspects were a security aspect responsible for generating and marshalling signatures in the client

and an aspect responsible for unmarshalling and checking signatures in the server.

In order to remove this crosscutting of the requirements, the aspects were factored as follows:

- A_G Aspect to generate a digital signature
- A_C Aspect to check a digital signature
- A_M Aspect to ensure signature is marshalled
- A_U Aspect to ensure a signature is unmarshalled

This allows us to factor the aspects across the system in the following manner. Aspects A_G and A_M are used at the ‘signing’ server, and aspects A_U and A_C at the ‘checking’ server. This factoring also allowed us to remove a degree of crosscutting with the original requirements. Our aspects now divide into two forms: those concerned with message manipulation (a low level mechanism) and those concerned with details specific to digital signatures (i.e. operating at more of a policy level). This distinction is significant when considering system evolution, as we discuss in section 4.

3.2 Aspect Details

Here we present the aspects that were developed to implement the digital signatures for our authentication service. In this discussion, we leave aside the issue of the distribution of the keys, and concentrate on some of the problems we encountered during this development.

On the *client* side, aspects A_M and A_G were implemented as shown in Figures 3 and 4 respectively. We defined a new class DigSig with a number of methods to create and check digital signatures and a class Packaging with methods to pack and unpack messages (these classes are not shown here). Aspect A_M defines a **pointcut** which intercepts the moment that messages are about to be sent to the server. The banking system communicates between its components using sockets, and the messages are sent as Strings using the method println(). The **around** advice ensures that marshalling takes place, using the pack() method from DigSig.

```
public aspect AM {

    pointcut marshal (String anArg): call(* *.println(..)
        && (args(anArg))
        && target(java.io.PrintWriter+);

    void around (String oldArg): marshal(oldArg) {
        String newArg = Packaging.pack (oldArg, null);
        proceed (newArg);
    }
}
```

Figure 3 Aspect A_M

A separate aspect A_G has its **pointcut** defined on the call to pack() within the **around** advice of the aspect A_M. This ensures that an

appropriate signature will be generated and passed on as a parameter to the pack() method of aspect A_M.

```
public aspect AG {

    pointcut generateSig (String message, String signature):
        within (AM)
        && call(* Packaging.pack(..)
        && (args(message, signature));

    String around (String message,String signature):
        generateSig (message,signature){
            return (proceed (message,(DigSig.CreateDigSig(message))));
        }
}
```

Figure 4 Aspect A_G

The reason for the approach taken here – in AM to pack with a null argument that is subsequently replaced by a digital signature by AG (createDigSig(message)) – is to enable this security mechanism to be changed in the future.

```
1 public void run() {
2     open client sockets for communication
3     while (true) //serve for duration of client session
4         try {
5             read a client request message
6             if (message is not in expected format){
7                 send Error message to client
8                 continue
9             } //end if
10            if (request is a login command) {
11                try{
12                    forward client details to proxy server
13                    get response from proxy server
14                    if (response == succesfull login){
15                        log client account number for this session
16                    }
17                    forward response to client
18                } catch (Exception e1)
19            } // end if
20            else
21                add message to queue
22        } catch (Exception e2){break}
23    } // end while
24    close sockets
25 } //end run
```

Figure 5 Pseudo code for the run method in QueueServerThread

On the *server* side it is the Queue Server where messages from the clients arrive and which must implement the checking activity. Within the Queue Server code is a class named `QueueServerThread` which extends Java's `Thread` class. In Java, the code that a `Thread` object executes must be implemented in an argument-less public void method named `run()`. The basic design of the `run()` method in `QueueServerThread` is shown in Figure 5. There is one thread for each client session.

The `run()` method is called from the `start()` method invoked by the constructor of `QueueServerThread`. The `run()` method contains a number of **try-catch** statements, as well as loop control statements, **break** and **continue**, which make it difficult to follow the flow of the program. This clearly illustrates why it would be best not to tamper with this application!

The point at which a client request enters the Queue Server is found on line 5. This represents a good join point to intercept the incoming message. Immediately the message arrives it is validated (lines 6 to 9) for things like missing command, missing password, or attempts to use services without authorization. If the message is found to be invalid, a message is sent to the client and the processing of the message stops – the thread returns to the beginning of the while loop to obtain the next client message. Note that currently the validation of the message does not include message tampering checks, and that is what the aspects will address by intercepting the message before it undergoes all its other validation checks.

Figure 6 shows the aspect A_U which takes care of the unmarshalling. It uses **around** advice, which first calls `proceed()` in order to obtain the result of the `readLine()` statement, it then converts this using the `unpack()` method from `DigSig`, and returns the modified message to the program.

```
public aspect AU{

    pointcut unmarshal() : call (* *.readLine(..))
                               target(java.io.BufferedReader+);

    String around () : unmarshal() {
        String signedMessage = proceed ();
        String unsignedMessage =
            Packaging.unpack(signedMessage);
        return (unsignedMessage);
    }
}
```

Figure 6 Aspect A_U

Finally, there is an aspect to check whether the digital signature is correct. Ideally, aspect A_C should be defined as follows: *'if during the run() method it is discovered that there is an intruder, then stop execution of the run() method for that client session'*. The idea being that the discovery of an intruder is sufficiently serious that we would not want it to continue processing, and contacting other servers with potentially harmful data.

We resolved to define aspect A_C as presented in Figure 7.

```
public aspect AC {

    pointcut checkSig (String uncheckedMessage):
        within(AU)
        && call(* Packaging.unpack(..))
        && (args (uncheckedMessage));

    before (String uncheckedMessage):
        checkSig(uncheckedMessage) {
        if (!DigSig.IsCorrectSig(uncheckedMessage)) {
            printWriter.println("E*Aspect intruder");
            //any other code to raise alarm
            throw (new SecurityException());
        }
    }
}
```

Figure 7¹ Aspect A_C

Notice that we defined a **pointcut** for A_C in the **around** advice of the aspect A_U as it does the unmarshalling, similar to the way aspect A_G was constructed on top of aspect A_M . Aspect A_C works as follows: it retrieves the unchecked message and ensures that it is checked through the `DigSig` class. If the result of this test is false, a message is sent to the client indicating that an intruder has been intercepted, and a `SecurityException` is thrown. It is this throwing of the exception which will ensure that the program will effectively abandon the `run()` method. This is because the exception thrown in the aspect will be caught on line 22 of the `run()` method which has the call to **break** to handle the exception.

It is worth noting that the design of the aspects conforms to a more general situation where messages are transmitted asynchronously between nodes and are required to be validated on receipt.

3.3. Altering the flow of program and the exception

In aspect A_C we deliberately allow the aspect to throw an exception, in order to cause the program to stop executing the method it was in (here, `run()`). Throwing an exception in an aspect, thus causing the original program to change its course in such a dramatic way, appears drastic. Before we resorted to this action a number of different strategies were explored:

- Rewrite the `run()` method entirely, and make sure that the new version is called instead of the original one. This is not possible, because the call to `run` is hidden within the `start()` method of the library `Thread` class. In addition, the complexity of the existing `run()` method is such that the risk of failing to re-implement it correctly is high.
- Alter the values of local variables in the `run()` method, for example setting the message to **null**, and then picking this up

¹ We have not shown the details of how to obtain access to the local variable, `printWriter`, representing the output stream to the client.

when the format of the message is checked (see line 6 in Figure 5). However, this strategy would tie our aspect a lot closer to the actual application code, making it less maintainable and less re-usable.

- Handle the exception using advice for the run() method. This would involve some complications, without mitigating the control flow issues of our chosen approach.

Given that the nature of the aspects is to improve the level of security in our application, and in particular to remove the danger of tampering with messages going into the bank system, there is some justification for the more serious action we chose.

However, this proposal violates two aspect style rules originally identified by Kersten and Murphy [9]:

- do not allow exceptions introduced by a weave to percolate out from the weave, and
- before and after advice should not alter the pre- and post-conditions of a method.

These two rules together ensure that the application of an aspect to a class will not affect how the class fits into the existing class structure, nor will it modify the contracts between client and supplier methods in the existing class structure, which in turn makes understanding, debugging and testing of the application much easier.

In aspect A_C the exception is thrown deliberately to ‘percolate out from the weave’. Had the exception been caught in the aspect, the desired effect would have been lost.

A preferred option would have been to use a dedicated aspectException, that is, an exception which can only be thrown by aspects. Assuming that such an exception would be a subclass of Exception it could still be caught within the run() method. Such an exception would facilitate managed interactions between an aspect and the control-flow of legacy code.

With respect to the second style rule, it is an interesting question whether the pre- and post conditions of the run() method are affected by the exception thrown in the aspect. We argue that the pre-condition is not affected, because the design of the run() method is such that no assumptions are made about the validity of the message string it is designed to work with. The post-condition is in fact strengthened by the behavior of the aspect, because the original post-conditions of the method (i.e. those before introducing aspects) still hold, and an additional one has been introduced. This is permissible under subcontracting [12] and hence indicates that we still operate according to the principles of Design by Contract.

The above suggests that there should be some modification to the Kersten and Murphy aspect style rules, to cater for a structured ‘percolation’ that still adheres to the principles of Design by Contract.

4. System Evolution

In reality it may well be that some of the servers in the banking system, those towards the back-end, will be adequately protected by other mechanisms than our digital signatures, including physical isolation, private networks, and proprietary security protocols. Additionally some of the servers towards the client end might be to some degree outside the scope of the bank’s own

security domains, for example, networks of ATMs run by affiliated organizations. As businesses evolve in the context of internal restructuring, mergers, and loose federal arrangements with heterogeneous organizations (such as large supermarket chains), the security boundary we established in Section 3 may be required to shift across machine boundaries.




Figure 8: Banking application server structure with amended security boundary (now Proxy defines the boundary)

In Figure 8 we show one possible new machine boundary, where the Proxy Server now defines the point where clients interface to the banking system. The solution presented in section 3 is able to gracefully accommodate this change. The reason for this lies in the structuring of our aspects. The low level mechanism on which we have built security – aspects to marshal and unmarshal digital signatures – is independent of the aspects that generate and check digital signatures. This supports the moving of our boundary: Aspects A_M and A_G both move to the Queue Server, and aspects A_U and A_C move to the Proxy Server.

During the development of this scenario an intermediate state may be required to test the system. This intermediate state could look as follows: the Bank Client with aspects to marshal and generate (A_M and A_G); the Proxy Server with aspects for unmarshalling and checking (A_U and A_C); and the Queue Server in between the two with aspects for marshalling and unmarshalling as well as generating signatures (A_U , A_M and A_G).

Furthermore, the structuring of the aspects means we can flexibly accommodate changes in security policy such as new algorithms, changes in key length and so on. Such changes in policy would be encapsulated in A_G and A_C and no changes to A_M and A_U would be required.

Finally, in Section 3 it was noted that the definition of the aspects displays minimal coupling with the application code. In this paper we have shown how they can be used to add digital signatures to messages under scenarios with differing security boundaries. In addition, given the generality of the aspects, particularly those defined for marshalling and unmarshalling, they will be able to provide support to scenarios entirely outside the realm of security. Applications with significant IO functionality can benefit from being able to easily validate or otherwise modify and enhance the input and output arguments by using the aspects.

5. Conclusions and Future Work

The contribution of this paper is in showing how aspects can be used to evolve legacy code. We have shown that taking one step of evolution (digital signatures), through an aspects approach, enhances future evolutions in a number of dimensions (moving

the security boundary, changes to security policy and during the development phase). We have also shown how a server's control flow behaviour can be modified through the use of aspects, and that this does not necessarily break the two aspect style rules as defined in [9].

Since our aspects relate to a single overall concern – authorization – their relationship is one of cooperation, rather than conflict. In a more realistic scenario, system evolution would be likely to involve simultaneous work across a number of cross-cutting concerns with a high likelihood of conflicting requirements. An approach to resolving such conflicts is given in [14]. Our approach to partitioning the aspects relating to a concern should allow conflict resolution to operate at a finer granularity: for example, replacing our signing algorithm with a faster one can be done independently of the marshalling aspects.

In future work we plan to investigate implementing further additions of functionality to our system. The aim being to explore a wider range of interaction patterns between aspects and legacy code, and also between the aspects themselves. From this we would hope to be able to identify some analysis and design patterns to encapsulate our approach in a general way. This will allow us to give further consideration to the question of whether there should be some modification to the Kersten and Murphy aspect style rules, to cater for a structured 'percolation' that still adheres to the principles of Design by Contract.

An alternative to our use of a method with a null argument in aspect A_M that gets replaced in aspect A_G , would be to use the Factory pattern, allowing us to vary our approach to signing. Further work is needed to consider whether, in general, this would be a better approach.

This work is a part of a larger programme of work on security, and in particular on security requirements. In contrast to the work presented here, which concerns itself with securing an implementation, we are also working on specifying security requirements independently of implementation [8].

6. ACKNOWLEDGMENTS

We gratefully acknowledge the input of our colleagues Charles Haley and Bashar Nuseibeh in the Department of Computing, The Open University, for their thought provoking discussions. Thanks, too, to the anonymous referees of an earlier version of this paper for their constructive suggestions.

REFERENCES

- [1] Aaltonen T., Helin J., Katara M., Kellomaki P., Mikkonen. T. Coordinating Aspects and Objects. *Electronic Notes in Theoretical Computer Science*, **68** No 3, Elsevier Science, 2003.
- [2] AspectJ Team. The AspectJ Programming Guide. <http://eclipse.org/aspectj>
- [3] Blaze, Feigenbaum and Lacy. Decentralized Trust Management. *IEEE Symposium on Security and Privacy*, Oakland, CA. May 1996.
- [4] Coady Y., Kiczales G. Back to the future: A Retroactive Study of Aspect Evolution in Operating System Code. In *Proceedings of AOSD03*, Boston, MA, 2003.
- [5] De Win B., Piessens F., Joosen W., Verhanneman T. On the importance of the separation-of-concerns principle in secure software engineering. *Workshop on the Application of Engineering Principles to System Security Design*, WAEPSSD, Boston, MA, USA, November 6-8, 2002, Applied Computer Security Associates (ACSA).
- [6] De Win B., Joosen W. and Piessens F. AOSD & Security: a practical assessment. *Workshop on Software engineering Properties of Languages for Aspect Technologies (SPLAT03)*, 2003, 1-6.
- [7] Grundy J. Aspect-Oriented Requirements Engineering for Component-based software systems. In *Fourth IEEE International Symposium on Requirements Engineering (RE'99)*. Limerick, Ireland: IEEE computer Society Press, 7-11 Jun 1999
- [8] Haley C.B., Laney R.C., Moffett J.D., Nuseibeh B.A. Using Trust Assumptions in Security Requirements Engineering. *Second Internal iTrust Workshop on Trust Management in Dynamic Open Systems*, Imperial College, London UK, 1-17 Sep 2003.
- [9] Kersten M.A., Murphy G.C. Atlas: A Case Study in Building a Web-based Learning Environment using Aspect-oriented Programming. *Technical Report Number TR-99-04*, Department of Computer Science, University of British Columbia, Canada, 1999.
- [10] Kiczales G., Hilsdale E., Hugunin J., Kersten M., Palm J., Griswold W. G. Getting Started with AspectJ. *Communications of the ACM*. October 2001, Vol. 44, No.10, 59-65.
- [11] Laddad R. *AspectJ in Action*. Manning, Greenwich, 2003.
- [12] Meyer B. (1997) *Object-Oriented Software Construction*. Prentice Hall International [2nd edition, ISBN 0-13-629155-4]
- [13] Nuseibeh B.A., "Weaving Together Requirements and Architecture", *IEEE Computer* 34 (3):115-117, March 2001.
- [14] Rashid A., Moreira A.M.D., Araujo J. Modularisation and Composition of Aspectual Requirements. In *Proceedings of AOSD03 2003*, Boston, MA, 2003.
- [15] Viega J., Bloch J.T., Chandra P. Applying Aspect-Oriented Programming to Security. *Cutter IT Journal*, **14**, No.2, February 2001, 31-39.
- [16] Welch I.S. and Stroud R.J. Re-engineering Security as a Crosscutting Concern. *Computer J.*, **46** (5), 578-589, 2003.
- [17] Zhang C., Jacobsen H. Quantifying Aspects in Middleware Platform. In *Proceedings of AOSD03*, Boston, MA, 2003.