

Autonomic management policy specification in Tune

Laurent Broto¹, Daniel Hagimont², Patricia Stolf³, Noel Depalma⁴, Suzy Temate⁵

¹Université Paul Sabatier, broto@irit.fr

²Institut National Polytechnique de Toulouse, hagimont@enseeiht.fr

³Institut Universitaire de Formation des Maîtres, stolf@irit.fr

⁴Institut National Polytechnique de Grenoble, depalma@inrialpes.fr

⁵Université de Yaoundé I, suzytemate@gmail.com

ABSTRACT

Distributed software environments are increasingly complex and difficult to manage, as they integrate various legacy software with specific management interfaces. Moreover, the fact that management tasks are performed by humans leads to many configuration errors and low reactivity. This is particularly true in medium or large-scale distributed infrastructures.

To address this issue, we explore the design and implementation of an autonomic management system. The main principle is to wrap legacy software pieces in components in order to administrate a software infrastructure as a component architecture. However, we observed that the interfaces of a component model are too low-level and difficult to use. Therefore, we introduced higher-level formalisms for the specification of deployment and management policies. This paper overviews these specification facilities that are provided in the Tune autonomic management system.

Categories and Subject Descriptors

D.2.11 [Software Engineering]: Software Architectures—*component-based autonomic management*

Keywords

Administration, autonomic system, deployment, reconfiguration

1. INTRODUCTION

Today's computing environments are becoming increasingly sophisticated. They involve numerous complex software that cooperate in potentially large scale distributed environments. These software are developed with very heterogeneous programming models and their configuration facilities are generally proprietary. Therefore, the management¹ of these software (installation, configuration, tuning,

¹we also use the term administration to refer to management operations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'08 March 16-20, 2008, Fortaleza, Ceará, Brazil

Copyright 2008 ACM 978-1-59593-753-7/08/0003 ...\$5.00.

repair ...) is a much complex task which consumes a lot of human resources.

A very promising approach to this issue is to implement administration as an autonomic software. Such a software can be used to deploy and configure applications in a distributed environment. It can also monitor the environment and react to events such as failures or overloads and reconfigure applications accordingly and autonomously. Many works in this area have relied on a component model to provide such an autonomic system support [4], [5], [8]. The basic idea is to encapsulate the managed elements (legacy software) in software components and to administrate the environment as a component architecture. Then, the administrators can benefit from the essential features of the component model, encapsulation, deployment facilities and reconfiguration interfaces, in order to implement their autonomic management processes. In a previous project (Jade [5]), we designed and implemented such a component-based autonomic management system. In the Jade system, an administrator can wrap legacy software in components (Jade relies on the Fractal component model [2]), describe a software environment to deploy using the component model ADL (Architecture Description Language) and implement reconfiguration programs (autonomic managers) using the component model's interfaces (Java interfaces in Fractal).

However, we observed that the interfaces of a component model are too low-level and difficult to use. In order to implement wrappers (to encapsulate existing software), to describe deployed architectures and to implement reconfiguration programs, the administrator of the environment has to learn (yet) another framework, the Fractal component model in the case of Jade. Tune is an evolution of Jade which aims at providing a higher level formalism for all these tasks (wrapping, deployment, reconfiguration). The main motivation is to hide the details of the component model we rely on and to provide a more intuitive policy specification interface for wrapping, deployment and reconfiguration.

The rest of the paper is structured as follows. Section 2 presents the context of our work and our motivations. Section 3 describes our contributions. After an overview of related works in Section 4, we conclude in Section 5.

2. CONTEXT

In this section, we first present an application case that we use to illustrate our contributions. We then present in more details what we mean by component-based autonomic management.

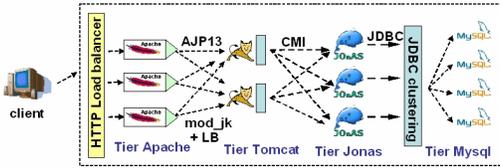


Figure 1: Clustered J2EE servers

2.1 Applications

Our main application target is the administration of servers distributed over a cluster of machines or a grid infrastructure. We give two examples of such organizations.

2.1.1 Clustered J2EE server

The Java 2 Platform, Enterprise Edition (J2EE) defines a model for developing web applications [7] in a multi-tiered architecture. Such applications are typically composed of a web server (e.g. Apache), an application server (e.g. Tomcat) and a database server (e.g. MySQL). Upon an HTTP client request, either the request targets a static web document, in which case the web server directly returns that document to the client; or the request refers to a dynamically generated document, in which case the web server forwards that request to the application server. When the application server receives a request, it runs one or more software components (e.g. Servlets, EJBs) that query a database through a JDBC driver (Java DataBase Connection driver). Finally, the resulting information is used to generate a web document on-the-fly that is returned to the web client.

In this context, the increasing number of Internet users has led to the need of highly scalable and highly available services. To face high loads and provide higher scalability of Internet services, a commonly used approach is the replication of servers in clusters. Such an approach (Figure 1) usually defines a particular software component in front of each set of replicated servers, which dynamically balances the load among the replicas. Here, different load balancing algorithms may be used, e.g. Random, Round-Robin, etc.

2.1.2 Distributed load balancer over a grid

Grid computing aims at enabling the sharing, selection, and aggregation of geographically distributed resources, dynamically at runtime depending on their availability, capability, cost and user's quality of service requirements. Diet [3] is an example of middleware environment which aims at balancing computation load over a grid. It is built on top of different tools which are able to locate an appropriate server depending on the client requested function, the data location (which can be anywhere on the system, because of previous computations) and the dynamic performance characteristics of the system. The aim of Diet is to provide transparent access to a pool of computational servers at a very large scale.

As illustrated in Figure 2, Diet mainly has the following components. A client is an application which uses DIET to solve problems. A Master Agents (MA) receive computation requests from clients. Then a MA chooses the best server and returns its reference to the client. The client then sends the computation request to that server. Local Agents (LA) aim at transmitting monitoring information between servers and MAs. LAs don't take scheduling decision, but al-

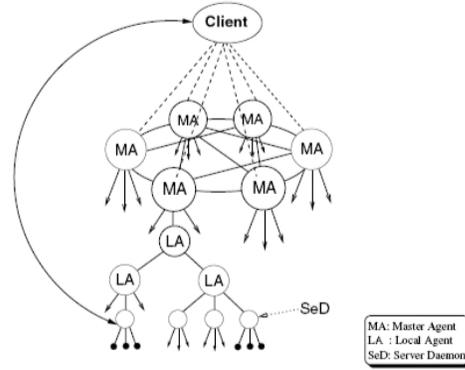


Figure 2: The Diet distributed load balancer

low preventing MAs overloads when dealing with large-scale infrastructures. Server Daemons (SeD) encapsulate computational servers (processors or clusters). A SeD declares the problems it can solve to its parent LA and provides an interface to clients for submitting their requests.

In both examples, the deployment of the servers in a cluster or grid is very complex and requires a lot of expertise. Many files have to be edited and configured consistently. Also, failures or load peaks (when the chosen degree of replication is too low) must be treated manually.

2.2 Component-based autonomic management

Component-based management aims at providing a uniform view of a software environment composed of different types of servers. Each managed server is encapsulated in a component and the software environment is abstracted as a component architecture. Therefore, deploying, configuring and reconfiguring the software environment is achieved by using the tools associated with the used component-based middleware. This solution is followed by several research projects, including Jade and Tune.

The component model we used in Tune is the Fractal component model [2]. A Fractal component is a run-time entity that is encapsulated and has one or more interfaces (access points to a component that supports a finite set of methods). Interfaces can be of two kinds: server interfaces, which correspond to access points accepting incoming method calls, and client interfaces, which correspond to access points supporting outgoing method calls. The signatures of both kinds of interface can be described by a standard Java interface declaration, with an additional role indication (server or client). Components can be assembled to form a component architecture by binding components interfaces (different types of bindings exists, including local bindings and distributed RMI-like bindings). An Architecture Description Language (XML based language) allows describing an architecture and an ADL launcher can be used to deploy such an architecture. Finally, Fractal provides a rich set of control interfaces for introspecting (observing) and reconfiguring a deployed architecture.

Any software managed with Tune is wrapped in a Fractal component which interfaces its administration procedures. Therefore, the Fractal component model is used to implement a management layer (Figure 3) on top of the legacy layer (composed of the actual managed software). In the management layer, all components provide a management

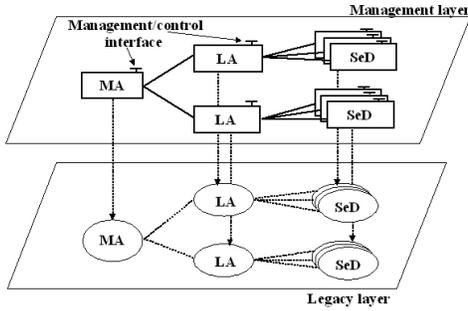


Figure 3: Management layer for the Diet application

interface for the encapsulated software, and the corresponding implementation (the wrapper) is specific to each software (e.g. the Apache web server in the case of J2EE or the Master Agent in the case of Diet). Fractal's control interfaces allow managing the element's attributes and bindings with other elements, and the management interface of each component allows controlling its internal configuration state. Relying on this management layer, sophisticated administration programs can be implemented, without having to deal with complex, proprietary configuration interfaces, which are hidden in the wrappers.

Here, we distinguish two important roles:

- the role of the management and control interfaces is to provide a means for configuring components and bindings between components. It includes methods for navigating in the component-based management layer or modifying it to implement reconfigurations.
- the role of the wrappers is to reflect changes in the management layer onto the legacy layer. The implementation of a wrapper for a specific software may also have to navigate in the component management layer, to access key attributes of the components and generate legacy software configuration files. For instance, the configuration of an Apache server requires to know the name and location of the Tomcat servers it is bound to.

2.3 Motivations

Component-based autonomic computing has proved to be a very convenient approach. The experiments we conducted with Jade [5] for managing J2EE or Diet infrastructures validated this design choice. But as Jade was used by external users (external to our group), we observed that:

- wrapping components are difficult to implement. The developer needs to have a good understanding of the component model we use (Fractal),
- deployment is not very easy. ADLs are generally very verbose and still require a good understanding of the underlying component model. Moreover, if we consider large scale software infrastructure such as those deployed over a grid (as in the Diet example), deploying a thousand of servers requires an ADL deployment description file of several thousands of lines,
- autonomic managers (reconfiguration policies) are difficult to implement as they have to be programmed

using the management and control interfaces of the management layer. This also required a strong expertise regarding the used component model.

All these observations led us to the conclusion that a higher level interface was required for describing the encapsulation of software in components, the deployment of a software environment potentially in large scale and the reconfiguration policies to be applied autonomically. We present our proposal in the next section.

3. TUNE'S MANAGEMENT INTERFACE

As previously motivated, our goal is to provide a high level interface for the description of the application to wrap, deploy and reconfigure. This led us to the following design choices:

- Regarding wrapping, our approach is to introduce a Wrapping Description Language which is used to specify the behavior of wrappers. A WDL specification is interpreted by a generic wrapper Fractal component, the specification and the interpreter implementing an equivalent wrapper. Therefore, an administrator doesn't have to program any implementation of Fractal component.
- Regarding deployment, our approach is to introduce a UML profile for graphically describing deployment schemas. First, a UML based graphical description of such a schema is much more intuitive than an ADL specification, as it doesn't require expertise of the underlying component model. Second, the introduced deployment schema is more abstract than the previous ADL specification, as it describes the general organisation of the deployment (types of software to deploy, interconnection pattern) in intension, instead of describing in extension all the software instances that have to be deployed. This is particularly interesting for applications like Diet where thousands of servers have to be deployed.
- Regarding reconfiguration, our approach is to introduce a UML profile for the description of state diagrams. These state diagrams are used to defined workflows of operations that have to be performed for reconfiguring the managed environment. One of the main advantage of this approach, besides simplicity, is that state diagrams manipulate the entities described in the deployment schema and reconfigurations can only produce an (concrete) architecture which conforms with the abstract schema, thus enforcing reconfiguration correctness.

We details these three aspects in the next sub-sections.

3.1 A UML profile for deployment schemas

The UML profile we introduce for specifying deployment schemas is illustrated in Figure 4 where a deployment schema is defined for a Diet organization. A deployment schema describes the overall organization of a software infrastructure to be deployed. At deployment time, the schema is interpreted to deploy a component architecture. Each element (the boxes) corresponds to a software which can be instantiated in several component replicas. A link between

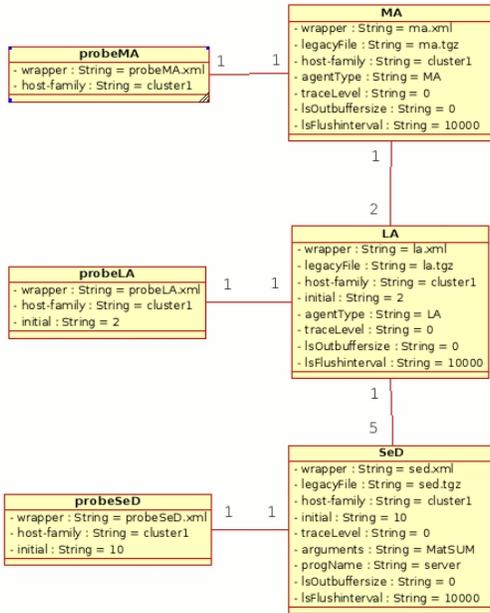


Figure 4: Deployment schema for Diet

two elements generates bindings between the components instantiated from these elements. Each binding between two components is bi-directional (actually implemented by 2 bindings in opposite directions), which allows navigation in the component architecture.

An element includes a set of configuration attributes for the software (all of type String). Most of these attributes are specific to the software, but few attributes are predefined by Tune and used for deployment:

- **wrapper** is an attribute which gives the name of the WDL description of the wrapper,
- **legacyFile** is an attribute which gives the archive which contains the legacy software binaries and configuration files,
- **hostFamily** is an attribute which gives a hint regarding the dynamic allocation of the nodes where the software should be deployed,
- **initial** is an attribute which gives the number of instances which should be deployed. The default value is 1.

The schema in Figure 4 describes a Diet organization where one MA, two LAs and 10 SeDs (5 for each LA) should be deployed. A probe is linked with each software, which monitors the liveness of the server in order to trigger a repair procedure. In this schema, a cardinality is associated with each link. If A(n) and B(m) are two linked elements in a schema, with an initial attribute (initial number of instances) n for A and m for B, the semantic of the cardinality is the following. A link A(n) t-u B(m) means that each A component should be bound with u B components and each B component should be bound with t A components. The cardinality is constrained by $m = n * u / t$ with $m \geq u$ and $n \geq t$.

The schema in Figure 4 deploys a component architecture as illustrated in Figure 3.

3.2 A Wrapping Description Language

Upon deployment, the above schema is parsed and for each element, a number of Fractal components are created. These Fractal components implement the wrappers for the deployed software, which provide control over the software. Each wrapper Fractal component is an instance of a generic wrapper which is actually an interpreter of a WDL specification.

A WDL description defines a set of methods that can be invoked to configure or reconfigure the wrapped software. The workflow of methods that have to be invoked in order to configure and reconfigure the overall software environment is defined thanks to an interface introduced in Section 3.3.

Generally, a WDL specification provides *start* and *stop* operations for controlling the activity of the software, and a configure operation for reflecting the values of the attributes (defined in the UML deployment schema) in the configuration files of the software. Notice that the values of these attributes can be modified dynamically. Other operations can be defined according to the specific management requirements of the wrapped software, these methods being implemented in Java.

The main motivation for the introduction of WDL are:

- to hide the complexity of the underlying component model (Fractal),
- that most of the needs should be met with a finite set of generic methods (that can be therefore reused).

Figure 5 shows an example of WDL specification which wraps a SeD computing server in a Diet architecture. It defines *start* and *stop* methods which can be invoked to launch/stop the deployed SeD software, and a *configure* method which reflects configuration attributes in the configuration file of the SeD software. The Java implementations of these methods are generic and have been used in the wrappers of most of the software we wrapped (LA, MA for Diet, but also Apache, Tomcat and MySQL for J2EE. We only had to add an implementation of a *configure* method for XML configuration files). A method definition includes the description of the parameters that should be passed when the method is invoked. These parameters may be String constants, attribute values or combination of both (String expressions). All the attributes defined in the deployment schema can be used to pass the configured attributes as parameters of the method invocations. However, several additional attributes are automatically added and managed by Tune:

- *dirLocal* is the directory where the software is actually deployed on the target machine,
- *compName* is a unique name associated with the deployed component instance,
- *PID* is the process identifier of the process that runs the software.

In Figure 5, the *start* method takes as parameters the shell command that launch the server, and the environment variables that should be set:

- *\$dirLocal/\$progName* is the name of the binary to be launched,
- *\$dirLocal/\$compName-cfg* is the name of the configuration file which is passed to the binary and which

```

<?xml version='1.0' encoding='ISO-8859-1' ?>
<wrapper name='sed'>
  <method name='start' key='appli.wrapper.util.GenericStart' method='start_with_pid_linux' >
    <param value='$dirLocal/$progName $dirLocal/$compName-cfg $arguments' />
    <param value='LD_LIBRARY_PATH=$dirLocal' />
  </method>

  <method name='configure' key='appli.wrapper.util.ConfiguresPlainText' method='configure' >
    <param value='$dirLocal/$compName-cfg' />
    <param value='=' />
    <param value='traceLevel.$traceLevel' />
    <param value='parentName.$LA.compName' />
    <param value='name.$compName' />
    <param value='lsOutbufferSize.$lsOutbufferSize' />
    <param value='lsFlushInterval.$lsFlushInterval' />
  </method>

  <method name='stop' key='appli.wrapper.util.GenericStop' method='stop_with_pid_linux' >
    <param value='$PID' />
  </method>
</wrapper>

```

Figure 5: A WDL specification

is generated by the `configure` method of the wrapper. `$arguments` is a parameter for the binary,

- `LD_LIBRARY_PATH=$dirLocal` is an environment variable to pass to the binary.

The `configure` method is implemented by the `ConfigurePlainText` Java class. This configuration method generates a configuration file composed of `<attribute,value>` pairs:

- `$dirLocal/$compName-cfg` is the name of the configuration file to generate,
- `=` is the separator between each attribute and value,
- and the attributes and value are separated by a `."` character.

It is sometimes necessary to navigate in the deployed component architecture in order to configure the software. For instance, in Diet, a LA has a configuration variable (in its configuration file) called `Name` which is a unique name associated with the launched server. This configuration variable is assigned with the `$compName` in its wrapper. A SeD which is a child of the LA must have a `parentName` configuration variable set to the name of the parent LA in its configuration file. Therefore, in the SeD wrapper (Figure 5), we need to access the `compName` of its parent LA in order to set this `parentName` configuration variable. Since in the deployment schema there is a link between the LA and SeD elements, there are bindings between the LA and the SeDs at the component level. These bindings allow navigating in the management layer. In Figure 5, the `parentName` configuration variable is assigned with the name of the LA component which the SeD is bound with.

3.3 A UML profile for (re)configuration procedures

Reconfigurations are triggered by events. An event can be generated by a specific monitoring component (e.g. probes in the deployment schema) or by a wrapped legacy software which already includes its own monitoring functions.

Whenever a wrapper component is instantiated, a communication pipe is created (typically a Unix pipe) that can be used by the wrapped legacy software to generate an event, following a specified syntax which allows for parameter passing. Notice that the use of pipes allows any software (implemented in any language environment such as Java or C++)

to generate events. An event generated in the pipe associated with the wrapper is transmitted to the administration node where it can trigger the execution of reconfiguration programs (in our current prototype, the administration code, which initiates deployment and reconfiguration, is executed on one administration node, while the administrated software is managed on distributed hosts). An event is defined as an event type, the name of the component which generated the event and eventually an argument (all of type String).

For the definition of reactions to events, we introduced a UML profile which allows specifying reconfiguration as state diagrams. Such a state diagram defines the workflow of operations that must be applied in reaction to an event. An operation in a state diagram can assign an attribute or a set of attributes of components, or invokes a method or a set of methods of components. To designate the components on which the operations should be performed, the syntax of the operations in the state diagrams allows navigation in the component architecture, similarly to the wrapping language.

For example, let's consider the diagram in Figure 6 (on the left) which is the reaction to a LA (software) failure in Diet. The event (`fixLA`) is generated by a `probeLA` component instance, therefore the `this` variable is the name of this `probeLA` component instance. Then:

- `this.stop` will invoke the `stop` method on the probing component (to prevent the generation of multiple events),
- `this.LA.start` will invoke the `start` method on the LA component instance which is linked with the probe. This is the actual repair of the faulting LA server,
- `this.LA.SeD.stop` will invoke the `stop` method on all the SeD component instances which are linked with this LA. This is necessary as in Diet, a restart of a LA requires to restart all its SeD children in order to reconnect to the LA. Here, the probes associated with the SeDs will trigger the restart of the SeDs,
- `this.start` will restart the probe associated with the LA.

Notice that state diagram's operations are expressed using the elements defined in the deployment schema, and are applied on the actually deployed component architecture. We are currently extending Tune to provide operations which re-deploy components (change location or add component instances) while enforcing the defined abstract deployment schema.

A similar diagram is used to start the deployed Diet environment, as illustrated in Figure 6 (on the right). In this diagram, when an expression starts with the name of an element in the deployment schema (LA or SeD ...), the semantic is to consider all the instances of the element, which may result in multiple method invocations. The starting diagram ensures that (1) configuration files must be generated, then (2) the servers must be started following the order MA, LA and SeDs. For each type of server, the server is started before its probe.

4. RELATED WORK

Autonomic computing is an appealing approach that aims at simplifying the hard task of system management, thus

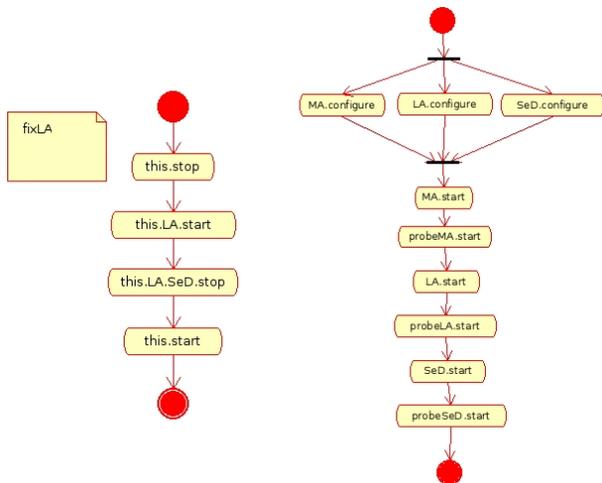


Figure 6: State diagrams for repair and start

building self-healing, self-tuning, and self-configuring systems [6].

Management solutions for legacy systems are usually proposed as ad-hoc solutions that are tied to particular legacy system implementations (e.g. [9] for self-tuning cluster environments). This unfortunately reduces reusability and requires autonomic management procedures to be reimplemented each time a legacy system is taken into account in a particular context. Moreover, the architecture of managed systems is often very complex (e.g. multi-tier architectures), which requires advanced support for its management.

Relying on a component model for managing legacy software infrastructure has been proposed by several projects [1], [4], [5], [8] and has proved to be a very convenient approach, but in most of the cases, the autonomic policies have to be programmed using the programming interface of the underlying component model (a framework for implementing wrappers, configuration APIs or deployment ADLs) which is too low level and still error prone.

In this paper, we proposed a high level interface which is composed of a language/framework for the description of wrappers:

- a UML profile for specifying deployment schemas,
- a UML profile for specifying reconfigurations as state transition charts.

The main benefit of this approach is to provide a higher level interface to the software environment administrator.

5. CONCLUSION

Distributed software environments are increasingly complex and difficult to manage, and their administration consumes a lot of human resources. To address this issue, many research projects proposed to implement administration as an autonomic software, and to rely on a component model to benefit from introspection and reconfiguration facilities that are inherent to these models.

Although component-based autonomic computing has proved to be a very convenient approach, we observed that the interfaces of a component model are too low-level and difficult to use. In order to implement wrappers for legacy software,

to describe deployed architectures and to implement reconfiguration programs, the administrator of the environment has to learn (yet) another framework with complex APIs or specific languages.

In this paper, we propose a higher level interface for describing the encapsulation of software in components, the deployment of a software environment and the reconfiguration policies to be applied autonomically. This management interface is mainly based on UML profiles for the description of deployment schemas and the description of reconfiguration state diagrams. A tool for the description of wrapper is also introduced to hide the details of the underlying component model.

We are currently extending our prototype to enable various form of reconfigurations (not just invocations on wrapper's methods). Notably, we provide support in state diagrams to allow component re-deployment, i.e. changing a component's location and adding a component instance, while still conforming to the specified abstract deployment schema.

6. ACKNOWLEDGMENTS

The work reported in this paper benefited from the support of the French National Research Agency through projects Selfware (ANR-05-RNTL-01803), Scorware(ANR-06-TLOG-017) and Lego (ANR-CICG05-11).

7. REFERENCES

- [1] G. Blair, G. Coulson, L. Blair, H. Duran-Limon, P. Grace, R. Moreira, and N. Parlavantzas. Reflection, self-awareness and self-healing in openorb. In *1st Workshop on Self-Healing Systems, WOSS 2002*, 2002.
- [2] E. Bruneton, T. Coupaye, M. Leclercq, V. QuÃ©ma, and J.-B. Stefani. The fractal component model and its support in java. In *Software - Practice and Experience, special issue on "Experiences with Auto-adaptive and Reconfigurable Systems"*, 36(11-12):1257-1284, September 2006.
- [3] P. Combes, F. Lombard, M. Quinson, and F. Suter. A scalable approach to network enabled servers. In *7th Asian Computing Science Conference*, January 2002.
- [4] D. Garlan, S. Cheng, A. Huang, B. Schmerl, and P. Steenkiste. Rainbow: Architecture-based self adaptation with reusable infrastructure. In *IEEE Computer*, 37(10), 2004.
- [5] D. Hagimont, S. Bouchenak, N. D. Palma, and C. Taton. Autonomic management of clustered applications. In *IEEE International Conference on Cluster Computing*, 2006.
- [6] J. O. Kephart and D. M. Chess. The vision of autonomic computing. In *IEEE Computer Magazine*, 36(1), 2003.
- [7] S. Microsystems. Java 2 platform enterprise edition (j2ee).
- [8] P. Oriezy, M. Gorlick, R. Taylor, G. Johnson, N. Medvidovic, A. Quilici, D. Rosenblum, and A. Wolf. An architecture-based approach to self-adaptive software. In *IEEE Intelligent Systems* 14(3), 1999.
- [9] B. Urgaonkar, P. Shenoy, A. Chandra, and P. Goyal. Dynamic provisioning of multi-tier internet applications. In *2nd International Conference on Autonomic Computing (ICAC-2005)*, Seattle, WA, June 2005.