A Comprehensive Solution for Application-Level Adaptation

K. Geihs^{1,*}, P. Barone², F. Eliassen^{3,7}, J. Floch⁴, R. Fricke⁵, E. Gjørven³,
S. Hallsteinsen⁴, G. Horn⁴, M. U. Khan¹, A. Mamelli², G. A. Papadopoulos⁶,
N. Paspallis⁶, R. Reichle¹, E. Stav⁴

¹ University of Kassel, D-34121 Kassel, Germany

² HP Italy Innovation Center, I-20063 Cernusco sul Naviglio (Milano), Italy

³ Simula Research Laboratory, NO-1325 Lysaker, Norway

⁴ SINTEF ICT, NO-7465 Trondheim, Norway

⁵ Condat AG, D-10559 Berlin, Germany

⁶ University of Cyprus, CY-1678 Nicosia, Cyprus

⁷ University of Oslo, NO-0316 Oslo, Norway

* Corresponding author: geihs@uni-kassel.de

ABSTRACT

Driven by the emergence of mobile and pervasive computing there is a growing demand for context-aware software systems that can dynamically adapt to their run-time environment. We present the results of project MADAM which has delivered a comprehensive solution for the development and operation of context-aware, self-adaptive applications. The main contributions of MADAM are (a) a sophisticated middleware that supports the dynamic adaptation of component-based applications, and (b) an innovative model-driven development methodology which is based on abstract adaptation models and corresponding model-to-code transformations. MADAM has demonstrated the viability of a general, integrated approach to application-level adaptation. We discuss our experiences with two real-world case studies that were built using the MADAM approach.

KEYWORDS: adaptive software, component framework, compositional adaptation, context awareness, middleware, model-driven development, utility function

1 Introduction and Motivation

Mobile and pervasive computing introduce a growing demand for software systems that are able to adapt to dynamically changing environments. Generally, we differentiate between static and dynamic adaptation: Static adaptation relates to the redesign and reconfiguration of application architectures and components at design time when functional and non-functional requirements change. Dynamic adaptation happens at application run-time due to changing resource and context conditions. For example, applications may want to react dynamically to fluctuations in network connectivity, battery capacity, appearance of new devices and services, or to a change of user preferences.

In this paper we focus on dynamic adaptation of context-aware applications on mobile computing devices. We present the results of MADAM, a joint European research project with eight partners from industry and academia from all over Europe [1]. Mobile computing has been chosen as our main application scenario since dynamic changes of the execution context and the need for application adaptation will be evident especially in such scenarios where users move from one location to another, e.g. from home to office to customer premises.

Clearly, MADAM has not been the first project to tackle dynamically adaptive software systems. It is based on, incorporates, and extends earlier research work. However, MADAM stands out since it provides a truly comprehensive solution that addresses adaptation from both the theoretical and practical perspective and solves all of the following challenges:

- ◆ Adaptation happens seamlessly and without user intervention in reaction to context changes.
- ◆ Applications may exploit any kind of context dependencies as long as there is appropriate hardware and software available in the computing environment to provide this context data.

- ◆ Context-awareness and adaptivity of applications are treated as a separate concern in application design.
- ◆ A general component model and middleware infrastructure supports many adaptation styles, e.g. local and distributed adaptation, parameter and compositional adaptation, anticipated and unanticipated adaptation.
- ◆ A model-driven development approach comprising adaptation models and corresponding transformations facilitates the development of self-adaptive applications and the reuse of adaptation artefacts.
- ◆ Real applications from industry partners are used to evaluate the approach.

While many different aspects of adaptive software systems have been explored before (see Section 6 and the special issue [2]), in MADAM we have taken a generalised research approach that integrates the majority of those aspects. Thus, the main contributions of this paper are:

- ◆ We demonstrate the viability of a general, integrated approach to application-level adaptation for context-aware applications.
- ◆ We present a flexible system infrastructure that enables dynamic adaptation in many ways; it can be realised on top of different component platforms such as Java and .NET.
- ◆ We introduce a new general-purpose, model-based development methodology for adaptive applications that builds on UML and Eclipse.

In this paper we present a retrospective on the overall approach, main achievements and general conclusions of project MADAM. Due to space limitations and the large amount of information we cannot go into full detail in some sections. In these cases we refer the reader to previous publications where more details can be found. The structure of the paper is as follows. Section 2 discusses adaptation requirements and explains how these requirements are translated into a novel generalized adaptation approach. Section 3 presents a component-based computing infrastructure that supports dynamic adaptation at the application level. In Section 4 we address the question of how to assist the developer in the design and implementation of general-purpose adaptive applications. A model-driven development methodology is presented that eases considerably the development task. In Section 5 we describe two real-world case studies that were built using the MADAM approach, and we discuss our experiences and lessons learned. Related works are analyzed in Section 6. Finally, Section 7 presents conclusions and outlook to future work.

2 Adaptation Requirements and Approach

The increasing proliferation of mobile devices has affected the expectations of their users. Mobile users now expect their devices to automatically and intelligently adapt to their ever-changing context. However, such an adaptive behaviour can be quite difficult to develop, especially when the nature of the adaptations is complex and hard to describe. Adding to the complexity is the fact that different sets of context information are available on different devices and also that different users have varying preferences concerning the desired adaptations.

The work of the MADAM project has endeavoured at providing solutions which ease the task of developing context-aware, adaptive systems for mobile computing. As a first step in the project, we sketched a number of scenarios that were then used to extract a set of requirements in the form of reference adaptations. This list of requirements has guided us through the design of the MADAM middleware and methodology.

2.1 Requirements

The MADAM case studies were based on both real commercial products and on hypothetical scenarios. In the following we describe a simplified scenario which is derived from one of the case studies: the SatMotion application. This scenario is revisited in subsequent sections. For the complete

scenario and the case study, the interested reader is referred to the project deliverables D1.2 [3], D5.3 [4], and D6.2 [5].

2.1.1 *Application scenario*

Consider a maintenance worker who uses an application running on her mobile device (e.g. a PDA) for her on-site work. Because of the nature of her tasks, the worker might or might not be able to visually interact with the mobile device. Consequently, the application is designed to offer multiple modalities for user interaction, i.e. visual and audio interaction. The visual interaction displays text messages (e.g. in popup windows) while the audio interaction uses the PDA's speakers and microphone for interaction with the user. Depending on the user's context, the middleware automatically selects which of the available modes to choose. Additionally, the application is capable of running a speech-to-text service both locally (on the master node) and remotely (on a slave node). When the audio interaction is used, the adaptation reasoning mechanism automatically decides whether hosting the text-to-speech engine locally is more beneficial than hosting it remotely on a server, and vice versa.

For example, consider the case where the worker is in her office, interacting with the PDA in order to prepare herself for an on-site visit. In this case, the application has the full attention of the user, and consequently the visual interaction is selected as it is more responsive (e.g. faster interaction) as well as more resource efficient (i.e. requires less memory and CPU, and uses no networking). But when the worker moves to the site and starts working with the equipment, the application switches to audio interaction, thus releasing the worker from the burden of having to visually interact with the application. Finally, when the network is sufficiently fast and cheap, and the resources are low (e.g. because the PDA starts other applications as well), the application switches to the audio mode where the speech-to-text component is hosted on a remote predefined server.

2.1.2 *Reference requirements*

Based on scenarios such as the one above, we concluded a set of reference requirements in the form of possible adaptations that the MADAM solution should support:

User Interface delegation: Transfer of (part of) the UI functionality to another device. For example, while driving a car a mobile user may have the option of connecting the mobile device to the onboard car computer in order to enable a hands-free interaction mode with the mobile device.

User Interface presentation: Tune the UI such that it is optimized for the current context conditions. For example, based on the ambient light conditions of its location, a mobile device could adjust its screen brightness.

User and application session redeployment: Redeployment of components, and/or applications to different devices, to accommodate for more efficient use. For example, in the office parts of the user application can be redeployed from her handheld to her desktop PC.

Functional richness: Extension of the functionality of an application by providing access to new software or hardware components. For example, an application which controls the home cinema in a living room could extend its functionality if a new hardware plug-in was added for controlling the lights of the room.

Data richness: React to the dynamically changing quality of a data stream. For example, a video conference application might want to adapt to different network characteristics, such as available bandwidth and latency by adjusting the resolution and the colour-depth of the video.

Network availability: Adapt to different mobile network technologies, e.g. GPRS, WLAN, Bluetooth, etc. For example, a mobile device can switch from a WLAN connection when the user is in the office to a GPRS connection when the user is leaving the building.

Security: Adjust the security mode based on the execution context. For example, communication via a GPRS communication link requires strong encryption, but when the user returns to the office and connects to the local WLAN, encryption may be turned off, thus saving resources on the device.

Software mode: Switch between different application-specific modes of operations. For example, a video streaming system may select different compression algorithms depending on data rates and communication quality.

Data replication and synchronization: Automate the replication and synchronization of data both in space and in time. For example, switch from local data to a server data base when the communication link is fast and cheap, or defer the downloading of large documents until the network is sufficiently fast or cheap.

Changing user preferences: Cater for changing user needs and preferences during the application lifecycle. Thus, the user needs and preferences should be considered as part of the context. For example, the user may want to specify different application priorities for different situations, such as work and leisure.

Although these adaptation types appear to be sufficiently simple at an individual level, it is argued that programming a system to achieve several of these adaptations in parallel can be quite challenging. Additionally, it is rather obvious that the development process for adaptive applications involves a number of repeated steps, such as defining how to collect the context information and how to decide about the adaptations. Hence, there is a need for a more methodical and automated approach in order to rid the developers from complex and repetitive work.

2.1.3 *Adaptation constraints*

While the above requirements represent items on our wish list for a general adaptation framework, we have to take into account certain inherent constraints and limitations. The most important ones are related to the resource constraints on mobile devices and to the distributed nature of applications.

Typically, mobile computing devices are constrained in terms of processing power, storage capacity and communication capabilities. Furthermore, battery power is limited. Thus, resource consumption and scalability of adaptation infrastructures are major concerns. Generally, the adaptation overhead is a crucial point in the design of such systems. In order to achieve scalability the number of variants that are evaluated at run-time in reaction to a context change, must be kept as low as possible by trying to eliminate configurations upfront that are not feasible in the given situation. In Section 4 we will present the solution which was adopted by MADAM.

Distribution is inherent to mobile applications and an important aspect in the design of adaptation frameworks. Additionally, distribution is useful as a means of alleviating the resource constraints of mobile computing devices. Hence, we must acknowledge the fact that applications may be distributed over client and server nodes, and that application components running on clients may access communication, processing and storage services provided on a server. Consequently, application adaptation decisions may involve local and remote contexts and application components, i.e. MADAM provides distributed context and adaptation management.

2.2 **Adaptation Approach**

From the discussions so far, it is obvious that multiple context elements may be relevant in a mobile setting and several forms of adaptation may contribute to enhance or maintain the service quality when the context changes. This section deduces the design of a general architecture for self-adapting systems and explains the adaptation rationale.

Three main middleware functionalities are required for the adaptation of applications:

- ◆ Monitor the context, detect context changes, and reason about the relevance of these changes.
- ◆ Make decisions about what adaptation to perform.
- ◆ Implement the adaptation choices, i.e. reconfigure the application(s).

First, let us consider context monitoring. We observe that multiple context elements need to be taken into account. Further, these span from elementary elements, such as network cost, to more complex aggregated or derived elements, such as predicted location. We expect that the set of relevant context elements and the sources producing them will evolve over time in the same way as applications do. Our experience with defining mobile scenarios is that often new context elements are introduced gradually as the problem to be solved becomes better understood or new application functionalities are added. All these observations lead to the following architectural implications:

i) Context monitoring should be kept separate from the application and realized through reusable context middleware. The context middleware should be extensible and support the addition of new context elements and new forms of reasoning.

Beyond reusability and extensibility of context components, the separation of concerns enforced by these implications should simplify the development of adaptive applications. Context management becomes the responsibility of components outside the applications. The context middleware can be developed, modified and extended independently from the applications.

Second, concerning adaptation reasoning, we observe multiple relations between context and adaptation mechanisms as well as interfering adaptation effects. Thus, context elements cannot be considered separately when reasoning on adaptation. For example, the audio capabilities should be considered together with the user activity before selecting the UI modality (i.e. voice or text-based UI). Also the implementation of an adaptation may have effect on the context. For example, the selection of the UI modality has an impact on the consumption of resources of the handheld device.

The more context elements we introduce, the more relations we need to deal with. From our experience with the generalization of scenarios we know that it is rather difficult to capture all these relations [18]. We also expect that new relations will be introduced as applications and context elements evolve. In [19] two kinds of approaches have been examined for self-adapting, context-aware applications: internal approaches where context-management and adaptation are realized as part of the application using programming language features, and external approaches where these mechanisms are realized by an application-independent middleware. The main drawback of internal approaches is the complexity introduced by intertwining adaptation and application behaviours. Also, they poorly support the evolution of the software. These drawbacks make internal approaches inappropriate in the context of mobile services, and thus we chose externalization of adaptation mechanisms:

ii) Adaptation mechanisms should be realized externally to the application.

External approaches require adaptation policies to be described separately from the applications. These policies are used by the middleware to reason and decide about adaptation. Three main approaches have been proposed for the description of such policies:

- ◆ Situation-action approaches specify exactly what to do in certain situations [20, 21]. Action-based policies are undoubtedly the most popular and are used in different domains related to networks and distributed systems such as computer networks, active databases and expert systems.
- ◆ Goal-oriented approaches represent a higher-level form of behavioural specification that establishes performance objectives, leaving the system or the middleware to determine the actions required to achieve those objectives [22]. Goal specifications capture the relations between context and adaptation mechanisms in a concise way.
- ◆ Utility-based approaches extend goal-oriented approaches [23, 24]. Utility functions are extended goal policies that ascribe a real-value scalar desirability to system states (i.e. in our case, a state is an application variant). The middleware computes the utilities of variants and selects the variant with the highest utility.

Situation-action approaches require the explicit description of each situation. The more context elements, the more conditions and rules need to be specified. Furthermore, these approaches use simple rules that fail to catch some dependencies between adaptation and context. They are therefore not appropriate for the purpose of our work. Using goal-oriented approaches, it is possible to express policies at a higher level. However, goal-oriented approaches also have major drawbacks. They fail to catch dependencies between adaptations and goals, and conflicts between goals. For example, a goal towards the optimization of CPU usage can be achieved through the suspension of lower priority applications, but such suspension would have impact on a goal towards high service availability. Also they do not provide any mechanism to compare adaptation actions when several actions can be applied to achieve a goal. Utility function policies express the rationale of an adaptation decision in a precise way, and are therefore more appropriate than goal policies when adaptation triggers and effects interfere, or when goals are in conflict. In our case, utility function policies can be specified to express

the dependencies between the properties of the system variants, the context and the user needs, and thus we chose a utility-based approach to tackle the complexity of adaptation reasoning:

iii) Adaptation policies should be expressed using utility functions.

Although utility-based approaches allow precise decision making, they require developers to specify in detail the properties of the application variants. It has turned out that for complex applications this might be a difficult task. Methodological help and appropriate tools are needed in order to support this.

Finally, we need to decide what kind of adaptation mechanism should be employed. Two popular adaptation techniques are parameter adaptation and compositional adaptation [12, 25]. Parameter adaptation supports fine tuning of applications through the modification of program variables and deployment parameters, while compositional adaptation allows the modification of the application component structure and the replacement of components. Parameterization is an intuitive and effective way to implement variability, but is less powerful than compositional adaptation and may also raise scalability concerns if the possible parameter value ranges are very large. Thus, our focus is on compositional adaptation:

iv) Adaptable applications should be built on compositional variability combined with support for component parameterization.

In summary, MADAM has adopted an external adaptation approach where adaptation capabilities of applications are realized by an application-independent middleware. Adaptation is based on compositional adaptation with added support for parameter adaptation. We follow an architecture-centric approach and employ architecture models at runtime to allow generic middleware components to reason about and control adaptation.

Figure 1 summarizes the main ideas of the MADAM adaptation approach. The middleware realizes three main adaptation-related functions: context management, adaptation management and configuration management. To fulfil these functions, the adaptation middleware needs to know the application structure and constraints, as well as the various context and resource dependencies. This implies that the middleware works at runtime on an architectural model of the application that specifies all adaptation variants and constraints. Note that this design enables unanticipated adaptations where new components or services are evaluated at runtime that were not available at application design time.

Figure 1: MADAM approach at a glance

The MADAM approach introduces a new major task for the developer: the building of application architecture models needed by the adaptation middleware to perform adaptation reasoning. Therefore, we have developed a software development methodology based on the Model Driven Architecture

(MDA) paradigm. The application developer captures the application variability in platform-independent adaptation models. The adaptation models are transformed by transformation tools into a representation (i.e. Java code in our case), that the middleware can access and leverage at runtime for adaptation management. This methodology is explained in detail in Section 4.

3 Adaptation Middleware

This section presents the design of the component model and computing infrastructure that satisfy the above requirements.

3.1 Adaptation Based on Component Frameworks

MADAM uses component frameworks to design applications that are capable of being adapted by reconfiguration. The component framework defines the required roles and interfaces of components that are plugged into the framework and regulates the interaction between components.

In MADAM, a component framework describes a composition of component types. These types represent variation points at which various component implementations can be plugged in. Variability is achieved through the plug-in of component implementations whose externally observable behaviour conforms to the component type. A component plugged into a component framework may be an atomic component, or a composite component built as a component framework itself. In this way, an application can be assembled from a recursive structure of component frameworks.

Figure 2 shows how components relate to applications and how variability is achieved. An application is considered as a component type that can have different realizations. The realization details of a component type are described using plans. Components can be atomic as well composite. Consequently, there are two types of plans: Blueprint and Composition Plans. A Blueprint Plan describes an atomic component and contains just a reference to the class or the data structure that realizes the component. The Composition Plan describes the internal structure of a composite component by specifying the involved Component Types and the connections between them. Variation is obtained by describing a set of possible alternative realizations of a Component Type using Plans.

MADAM also supports variability through component parameterization [13]. Component behaviour may be adapted by setting component parameters to appropriate values. This supports fine grained adaptation and may be used to fine-tune component behaviour such as tuning the encoding of a video stream to fit the available bandwidth. See Section 4 for details.

Figure 2: Creating Application Variants

To discriminate between alternative component implementations MADAM components are annotated with properties. Properties are used to qualify the services offered or needed by components. For

example, properties can describe that a user interface component implementation supports a hands-free mode. Properties can also specify requirements to system resources such as required memory space or network bandwidth. Properties are functionally dependent on context elements. In that way we are able to represent the dependencies between component implementations and context.

As components may provide several services and collaborate with multiple components, we need to distinguish between service quality on each collaboration. To do so, we use the concept of port. Components have ports through which they interact with each other. Thus, a port represents the capability of a component to participate in a specific kind of interaction or service. Properties are attached to ports. Services needed from the underlying platform such as memory are represented by implicit component ports, and platform related properties attached to these ports. Similarly, services provided to the user are represented by implicit component ports. It is the component type that defines that set of properties and their associations with ports.

Associated with implementations, property predictor functions are used to predict the properties of the implementations in a given context. These can be implemented in various ways including as constants, for instance to express that a specific component implementation requires more than 500 Kb memory, or as functional expressions of other properties of the component itself, of collaborating components or of inner components in the case of composite components. For instance, the response time offered by a component may depend on the network bandwidth available to that component, or the memory needed by a component may depend on the memory needed by inner components.

The concepts of property and port are illustrated in the example in Figure 3. The user interacts with the application in a specific context, as defined by the shown property values.

Figure 3: An application in context

Utility represents a measure of how well an application configuration fits a given context. The utility is given as a function of the properties of the ports representing the interaction between the application and its context. The utility function is specified by the application developer. In the current MADAM prototype, utility functions compute the weighted sum of the differences between offered and needed properties where the weights reflect the priorities of user needs. Within this model the goal of adaptation management can be formulated as ensuring that at any time the running system variant is the one with highest utility that does not violate the resource constraints imposed by the environment.

Distributed applications are designed using connectors that abstract communication protocols between remote components. Figure 4 shows an example of applying the MADAM component model. It illustrates the component model of a simple client-server application. The *STapplImpl* component is a composite component describing the component framework implementing the top-level application type *STapp* (service technician application). Each role defined by this framework, *UI* (user interface), *Ctrl* (control) and *Db* (database) is mapped to alternative component implementations at run-time. *Db* has two implementations that are also built as component frameworks.

In practice, as the service technician progresses through her daily tasks, the context changes either because she moves on to a different site, or because the environment changes (i.e. different time of day, different light conditions, weaker wireless network signals, etc). As a consequence, the possible configurations are evaluated using the utility function, and the application is appropriately adapted. For instance, the UI role is dynamically resolved to be implemented by either the visual (normal) or the audio (hands-free) UI, depending on the user context (are her hands busy driving or installing an antenna, or not).

Figure 4: Component framework architecture for the service technician application

3.2 Middleware Architecture

Figure 5 illustrates the MADAM middleware architecture and its most important interfaces. As explained before in Section 2, the middleware performs three main adaptation-related functions: context management, adaptation management and configuration management. We will present these middleware services in more detail in the following sections. In addition, two more components play a vital role inside the middleware, i.e. the Resource Manager and the Middleware Core.

The Resource Manager provides information on the state of the resources. It can be queried for the state of a resource in a pull mode, or components can subscribe to resource state change notifications

in a push mode. Furthermore, it handles locking and reservation requests, and collaborates with remote Resource Managers in case of distributed adaptation decisions. The Resource Manager is primarily accessed by the Context Manager for monitoring context related to the device resources (e.g. memory, CPU, networking) and by the Adaptation Manager during the adaptation process.

The MADAM Middleware Core provides platform-independent interfaces for the management of applications, plans, components and component instances. It is used to instantiate, initialize, and remove component instances, to connect and disconnect instances, to create remote bindings to remote components, to install, uninstall, and locate plans, and more. The Core and the Adaptation Manager collaborate in order to launch and manage applications. When an application launch request is received by the Core, it is forwarded to the Adaptation Manager, which in turn uses Core interfaces to locate, instantiate, and bind the components participating in the application.

Figure 5: Middleware architecture

3.3 Context Management

One of the main goals of the MADAM approach was to increase the reusability and to ease the development of context-aware, adaptive applications. In this respect, the context middleware was defined in a way which enables the developers to treat the collection and management of context information as an independent, cross-cutting concern.

3.3.1 Context model

By context, we refer to “any information that can be used to characterize the situation of an entity; [where] an entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and the application themselves” [14]. In our approach, we have adopted this definition and further we have classified context into three basic categories, based on the type of information it abstracts: user, computing, and environmental context [15]. Although many works in the literature have focused on certain types of context data, such as for example location-related context, in our work we have opted for a general approach, allowing the developers to specify,

model and use any type of context data. In this respect, a simple and extensible model was defined, allowing the developers to specify arbitrary types of context entities and context values.

The basic concept of this model is the *ContextEntity*. *ContextEntities* can include additional (child) *ContextEntities*, thus forming a tree-like structure. The leaves of this tree are the context values, which include numerical or string-based representations for the abstracted context data. This model also allows for the specification of a limited and predefined set of metadata, such as a time-stamp, the ID of the source which provided the context data, the probability that this data is correct, and the user rating when it is available. These metadata are primarily aimed for the internal context management mechanism, but they can also be accessed by context clients. More detailed descriptions of the MADAM context model are available in [16].

3.3.2 Context manager

The basic philosophy of the MADAM context architecture revolves around two points: First, context clients have specific context needs which they publish to the *ContextManager*. Second, context providers offer specific types of context data, which are also published to the *ContextManager*. For this purpose, the *ContextManager* provides a context repository which acts as a central hub for the routing of context information among sensors, reasoners and consumers, as it is illustrated in Figure 6.

Figure 6: Local and Distributed Context Management

The lifecycle of context data consists of several phases: collection of context data (by sensors), local and remote distribution, caching and storing, processing and eventually consumption of context data. Sensors are used to populate the context repository with context data. Context clients, such as external applications and the MADAM middleware register with the *ContextManager* for any of the two types of context change notifications: synchronous and asynchronous. The former provides a request/reply type of direct access to the context repository, while the latter uses the publish/subscribe pattern in order to enable a more event-based asynchronous notification.

For the purposes of the context middleware an extensive plug-in mechanism was developed, allowing the developers to dynamically add special mechanisms that feed the main *ContextManager* with context data. Thus, developers can dynamically add new functionality to the middleware, e.g. by packaging relevant context sensors along with the applications that need them, and the middleware can automatically start and stop context sensors depending on the context needs of the currently deployed applications. The latter enables the system to better utilize the resources, which are particularly scarce in the case of mobile devices.

Besides pluggable context sensors, the MADAM architecture also defines pluggable context reasoners. The latter are specialisations of sensors providing the additional capability of accessing context data from the repository with the purpose of processing it and deriving more advanced context information.

Once collected, the context data is preserved locally in a repository. The amount of data that is stored in the repository is configurable. The default is to only store the latest value of each context type, but this can be changed to allow for storing historical values. The latter is particularly useful in the case of context predictors, i.e. specialized context reasoners which try to infer the value of a context element in the future, with a given probability, by assessing historical context values.

Since MADAM primarily aims at distributed applications, one of the most advanced features of the *ContextManager* is distributed context management across multiple devices over either infrastructure-based or ad-hoc networks. This is achieved in a way that is transparent to both users and developers. For this purpose, a specialized component is defined: the *MembershipManager*.

The *MembershipManager* achieves two main tasks: management of group memberships and communication of context information across remote instances of the MADAM middleware. Groups are dynamically formed and contain instances of the middleware that are within reach of each other and willing to share their context information. The latter allows any member of the group to unicast or multicast any type of context information to other members of the group.

The basic assumption of this approach is that the context nodes publish their context requirements, and at the same time they collect and maintain the context needs of the other nodes. The distribution mechanism exploits this, and extends it across distributed instances of the middleware. In this respect, it is possible for a context client that resides on one device, to request a context entity of a type T, and receive it through a sensor which is attached to a different device within the same membership group. To enable this type of functionality, the membership managers implement a specialized protocol which periodically broadcasts messages encoding the (locally unsatisfied) context needs of the applications running on the devices. Hence, the context management space becomes broader.

Of course the distribution of context information has several implications, most notably on security and privacy aspects. As context data may include sensitive, user-related information, the context model is appropriately extended so that it classifies context types in different categories, according to their applicability or desirability for distribution. For more details on the modelling aspects, as well as on the local and distributed context management architectures the interested reader is referred to [17].

3.4 Adaptation Management

Each component may have a number of different realizations that provide the same basic functionality, but differ in their extra-functional characteristics like resource requirements and context dependencies, as described in Section 3.1. The MADAM component model describes how variability can be achieved by plugging in different component implementations whose externally observable behaviour conforms to the component type, but whose properties, which are described using the property model, are different.

In the MADAM middleware the *AdaptationManager* is responsible for reasoning on the impact of context changes on the applications, and for adapting the set of running applications by planning and selecting the application variants and the device configuration that best fit the current context. It consists of two subcomponents, the *AdaptationCoordinator* and the *BuilderandPlanner*. The *AdaptationCoordinator* is responsible for coordinating the actions to take place when a context change occurs and for selecting variants for the set of concurrent applications. In order to select the fittest variant, the *AdaptationCoordinator* needs to determine all the possible variants for each application, and to evaluate them. The *BuilderandPlanner* is responsible for dynamically building descriptions of alternative application variants (represented as configuration templates). The *AdaptationCoordinator* maintains a *BuilderandPlanner* for each concurrent application, and uses the configuration templates to assess the utility of the variants in the current context and select the “best” variants.

The *AdaptationManager* manages applications throughout their life cycle:

- ◆ **Application launch:** When a client starts a new application, the *AdaptationManager* builds a model of the adaptive application by retrieving the plans that describe how the application's component types can be realized, recursively traversing the composition structure. In order to be able to manage the adaptation of the application, it registers to receive notification about context changes relevant for that application. When the *AdaptationManager* has determined the best set of application variants, the *Configurator* (explained below) is requested to instantiate this application configuration, and to configure the device using the device settings determined to best match the current context and the running set of applications.
- ◆ **Application run-time adaptation:** Adaptation is triggered by context information received by the *AdaptationManager* from the *ContextManager*. Depending on the context information, the *AdaptationManager* may decide to replace the running application variant with an alternative variant, and invoke the *Configurator* to perform reconfiguration of the applications and the device.
- ◆ **Application shutdown:** On application shutdown, the *AdaptationManager* performs the necessary cleanup for the terminating application, and find the best configuration for the remaining set of running applications.

In order to enable distributed configurations in which applications depend on remote connections spanning distributed MADAM nodes, MADAM provides a distributed adaptation management protocol based on centralized adaptation reasoning, and distributed configuration. The application's master node makes the adaptation decisions, and delegates configuration or reconfiguration operations to the slave nodes. The MADAM nodes are connected by remote bindings.

Clearly, centralized decision making is subject to scalability problems. However, this is not expected to be a problem for typical MADAM deployments as mobile devices generally run just a few concurrent applications at a time.

3.5 Configuration Management

The responsibility of the MADAM configuration management is to orchestrate the changes required to safely reconfigure from the current set of running application variants to the new set of application variants selected by the adaptation management. This includes starting and stopping of applications and their components, configuring their connections and parameters, and changing device settings.

The main component handling the configuration management in MADAM is the *Configurator*. It is responsible for coordinating the initial configuration and later reconfigurations of the application components and the device. The *AdaptationManager* and the *Configurator* are tightly coupled as they operate on a common information element: the configuration template. When reconfiguring an application, the *Configurator* carries out the adaptations decided by the *AdaptationManager* by applying the provided configuration templates, one for each application. These templates contain complete descriptions of the choices made by the *AdaptationManager*, including the composition for each application and their structured subcomponents, blueprints for atomic components, mapping to the nodes in the infrastructure, connectors to use between components, and parameter settings for the components. The *Configurator* uses the information in the configuration templates to create and remove component instances on the local and remote nodes, to connect and disconnect components through local or remote connectors, and to set the parameters of the components.

The *Configurator* maintains information about the current set of running application variants. When requested to perform a re-configuration, the *Configurator* compares the new configuration templates with the information about the currently running applications to derive the (minimal) sequence of steps required to achieve the new configuration.

Figure 7: Application component reconfiguration states

Before starting a reconfiguration, the *Configurator* needs to drive the applications to a safe state, e.g. to avoid invalid interactions between application components while component instances and connectors are being replaced. This requires the collaboration of the components constituting the application, and the middleware expects configurable application components to implement an interface for changing their configuration state. This interface defines methods for setting the current configuration state of the components, for retrieving the current configuration state, and for transferring serializable internal component state from one component to another when the implementation of a specific component role is replaced by another implementation or moved to another node during reconfiguration. During the reconfiguration, components provide feedback to the *Configurator* when they are finished initiating or suspending. Figure 7 shows a state machine diagram for the configuration states of configurable components. The triggers for the state transfers are methods of the configuration interface of the application components and the feedback interface of the *Configurator*.

3.6 Implementation and Prototype

The MADAM project team identified a set of guidelines for the selection of the technical implementation platform. The guidelines have taken into account many factors, such as cross-platform portability, availability of IDE tools, availability of documentation and technical support for the candidate development tools and environments, considerations on performance and pricing, and leveraging existing component-based frameworks and patterns widely adopted by a well-established developers community. Since cross-platform portability is one of the main requirements for the MADAM middleware, the reference technologies selected for the development are Java-based, as the Java language facilitates the task of creating platform-independent solutions. However, it was not possible to make all the components completely independent from the platform, since the middleware needs to access low-level physical resources which requires some native code. We will come back to portability issues in more detail later in this section.

The resulting MADAM prototype can be deployed and executed on:

- ◆ Computing devices that offer a Java run-time environment supporting the J2ME CDC Personal Profile specification. Examples of such kind of environments, where the middleware has been tested, are PDAs running the MS Windows Mobile for Pocket PC operating system, enabled to run Java applications by means of the NsiCom CrE-ME 4.0 Java environment for PDA.
- ◆ Computing devices that offer a Java run-time environment supporting the Sun JRE J2SE 1.3.1 specification. Examples of such kind of environments, where the middleware has been tested, are computers running the MS Windows XP operating system, enabled to run Java applications by means of the official SUN Microsystems Java Running Environment 1.3.1 for Windows XP.

After an evaluation phase, the team selected as primary target devices some PDAs of the HP iPAQ family, like the iPAQ 6340 and 6915 running Microsoft Windows Mobile for Pocket PC. A second target category of devices is formed by Microsoft Windows XP-based laptops and tablet PCs. The selected Java environments are CrE-ME 4.0 from NsiCom Ltd for PDA, which is based on the JDK 1.3.1 specification, and the official SUN Microsystems Java 1.3.1 implementation for Windows XP in order to maintain code compatibility with the mobile Java platform. The development platform is MS Embedded Visual C++ 4.0 plus the MS Pocket PC SDK. The communication between embedded and Java components is based on JNI technology.

Figure 8 Examples of two MADAM instances running in the same adaptation domain

Figure 8 shows snapshots of two instances of the MADAM middleware running on two HP iPAQ 6340 PDA and sharing context data. In the following, some development related topics will be described in more detail.

Coding issues

The code for the MADAM middleware was developed by implementing for each component the interfaces as defined in the architectural design. The middleware developers applied a set of well-known and widely adopted design patterns that allow facing implementation issues by means of tested, proven development paradigms. For example, the Singleton pattern was employed for avoiding unwanted and unneeded proliferation of objects instances; the Observer pattern was used to implement subscription for listening to context change events and resource change events; the Adapter pattern allowed to wrap classes and methods from external or native APIs and to offer them in a format compliant with the component interface definition; the Proxy pattern was applied for providing a reference to remote objects.

In order to optimize the resource consumption and to provide better performance, some of the middleware components create and use Java threads, e.g. for monitoring low-level resources, handling the context distribution, performing adaptation steps, and handling GUI events. Since the thread usage may lead to a high consumption of system resources, which is a particularly critical issue on mobile devices, the number of threads has been kept limited and under the control of the MADAM middleware. Thus, they cannot dynamically proliferate in an unmanaged fashion.

The Context Manager runs a custom protocol for distributed context management. The protocol uses multicast messages (based on UDP over IP) for heartbeat signalling. It enables different instances of the middleware running in the same adaptation domain to communicate with each other and to share context data. Driven by the distributed context information provided by the Context Manager, the Adaptation Manager initiates distributed adaptations leveraging the Java RMI technology, which allows to invoke methods on remote MADAM instances. Similarly, the Resource Manager uses RMI to deal with resources of devices distributed over the network.

Network handling

The MADAM middleware prototype was designed to run in mobile environments, where the execution context changes dynamically. In such a scenario, it is fundamental to provide network access mechanisms that enable to automatically detect, select and use the best network connection among the available ones (if any). Therefore, the middleware relies on a pluggable mobile-IP API, whose tasks are handling transparently the network handover, providing information on the current state of the network connection and on the set of all available network types, and allowing connection to (disconnection from) a given network. The integration of the mobile-IP API into the MADAM

middleware takes place in the Resource Manager which provides generic methods for accessing network-related information. A NetworkAdapter class inside the Resource Manager acts as a wrapper for the pluggable mobile-IP API functionalities and provides them to the Resource Manager in a format compliant with its interfaces. The NetworkAdapter class is the only entry point in the whole middleware where external network APIs are accessed.

Cross-platform portability

The selection of a Java-based execution environment for the middleware allows a high degree of cross-platform portability. Nevertheless, as already mentioned, the implementation of some native code was unavoidable to collect information on low-level sensors and resources. The Resource Manager invokes such native code by leveraging the Java Native Interface technology. As a consequence, porting the middleware to platforms other than Windows XP or Windows Mobile requires the development of some custom native code for monitoring those low-level elements that are not directly accessible from Java. The points to modify in the code are anyhow well localized, enabling the developers to integrate new native code without the need for a complete knowledge of the middleware architecture.

An additional modification required for porting the middleware to other platforms is related to the above mentioned mobile-IP API. The middleware was designed to handle the external network API by means of the adapter class as introduced in the previous section. Hence, the usage of a different network API requires a reimplementation of the adapter class.

Lessons learned during development and testing

During the development and testing phases, the middleware development team faced some difficulties that were mainly related to the complexity of distributed environments and the limited resources of the mobile devices. Often these effects are quite subtle and hard to detect. Let us give two concrete examples.

In order to enable a distributed context data sharing among multiple MADAM nodes running on different devices, it was necessary to reach a compromise between the frequency of context data updates and their propagation over the network and the obvious impacts on the middleware performance. Updates in the context data, in fact, may trigger adaptations of the applications running on the middleware. The adaptation reasoning is a computing intensive task that, on constrained devices, may take full control of the device for a rather long time. When this happens, sometimes the middleware is neither able to receive heartbeat messages coming from remote Context Managers, nor able to generate new ones. As a consequence, it loses its membership to the current adaptation domain, and the whole process of adaptation fails.

Another area that required some extra fine tuning efforts was the resource management. The low-level resources are constantly monitored and some of them change their state at a high rate. If each such change in a resource would trigger a context change event, it would soon cause an overloading of the resource and context managers. To avoid this, each resource is checked according to predefined policies, e.g. at predefined intervals. Finding the appropriate values for such intervals without affecting negatively the behaviour and the adaptation capabilities of the applications running on top of the middleware required a lot of testing efforts.

4 Development Methodology

The development of context-aware, adaptive applications for dynamic operating environments is a challenging task. An important overall goal of MADAM is to facilitate the development of self-adaptive applications. We provide a model-based development methodology that builds on appropriate software modelling elements and automatic model transformations.

4.1 Adaptation Model

The adaptation model specifies the variability of the application including architectural constraints, component properties and context dependencies. We use UML 2.0 as modelling language. The standard UML 2.0 meta-model is enhanced by building a UML profile that provides the required modelling concepts for the adaptation capabilities. In order to explain our modelling approach in detail,

we demonstrate a step-by-step example using the SatMotion application, which was already introduced in Section 2. As in the previous sections, we show only a simplified version here.

The variability model starts with the specification of different possible realization options of the application at the top level. In our example the application can be realized through any of the three composition plans shown in Figure 9.

Figure 9: Different realization plans for the application

Each of the composition plans describes – among other things – a specific composition of component types that collectively realize the application. Figure 10 shows the composition of component types for the “Standard” composition plan of Figure 9.

Figure 10: A possible composition of components for the SatMotion application

The composition consists of four component types (marked with the `<<mComponentType>>` stereotype): Controller, MathProcessor, Recorder and UserInterface. Therefore, just like the application each of the component types may have a number of different realization plans. For example, the Controller component type has three different realizations (see Figure 11). A composition plan again consists of a composition of component types, while a blueprint plan describes the realization of a single component. Different variants of the application are created by resolving recursively all possible variation points; the recursion ends once a blueprint plan is reached.

Figure 11: Different possible realization plans for the Controller component type

Resolving all the variation points with all the possible options can effectively create a huge number of different application variants, all of which have to be evaluated at runtime in order to find the variant which is best for a given context condition. However, very often not all variants are actually feasible. For example, the selection of a particular realization for a component may imply a certain realization for another component. Likewise, a particular realization may be incompatible with some other component realizations. In order to ensure appropriate application variants for all context situations, infeasible variants have to be eliminated. Furthermore, the potential combinatorial explosion of variants can lead to a scalability problem requiring too much computational effort for a resource-scarce mobile device [6]. Therefore, we have added means to filter out the infeasible combinations. This can drastically reduce the number of variants to be considered at runtime. Our filtering approach builds on architectural constraints specified as features in the application model [7]. An example of a composition with architectural constraints is presented in Figure 12. The feature *CtrlMPIncompatibility* prevents the selection of incompatible components for *MathProcessor* and *Controller*; the feature *TypeMatching* enforces that only components with compatible types are selected for *Controller* and *UserInterface*.

Figure 12: Architectural constraints through feature specifications

The adaptation model also specifies the context and resource dependencies of the components, the component types and eventually the application. Both, context entities (and their types) and resources (and their types) are modelled hierarchically, including their offered and required properties. The context entities and resources are associated with the components and component types, and thus the context and resource dependency is established. An example resource model is presented in Figure 13. It shows resource types such as Memory, Network and Battery and their respective properties. Context entities are modelled the same way. For a more detailed discussion the reader is referred to [13].

Figure 13: Resource model

Adaptation in MADAM is built on compositional adaptation which enables the context-aware exchange and reconfiguration of application components at run-time. Thus, application algorithms and strategies can be modified by integrating components that may not have existed at application design time. Note that this approach supports unanticipated adaptation insofar as new component realizations may be added at runtime to the component framework, thus effectively enlarging the number of variants dynamically.

Compositional adaptation is a much more powerful technique than parameter adaptation where parameter values of components are adjusted as a reaction to context changes [12]. Nevertheless, we wanted to support parameter adaptation. Otherwise the developer might end up with a cumbersome explicit modelling of different parameter values as different component configurations. Therefore, we have integrated convenient modelling support for component parameters that may assume a range of (discrete) values. Internally, the middleware treats parameter adaptation just like compositional adaptation: The parameter model is transformed to different component compositions whose utility is evaluated by the adaptation manager separately.

In summary, the MADAM variability model captures all variation options as well as their utilities and context and resource dependencies in a concise, flexible and easily readable way. The platform-independent variability model is automatically transformed by a chain of tools into code that the MADAM middleware uses to perform adaptation decisions.

4.2 Model Transformation

In accordance with model-based development, the application adaptation model is built as a platform-independent model using an UML modelling tool. It is then transformed by means of a model-to-text transformation into code that performs adaptation-related functions such as evaluating context parameters and resource conditions. The adaptation middleware calls this code at runtime when resolving the variation points and evaluating the utility of possible variants. If there is a context change during application execution, the underlying adaptation middleware computes on-the-fly all possible

application variants and evaluates their fitness with respect to the current context situation. The “best” variant is selected and instantiated [8].

In MADAM we have used the Eclipse Modelling Framework, in which UML modelling tools such as Omondo and Borland Together Architect can be integrated. The UML application model is exported to XMI [10] in accordance with the meta-model defined by the Eclipse UML2, which is a lighter version of the OMG UML 2.0 specification. The UML2 model exported as XMI is taken as input to generate programming language code using MOFScript [9], which is an implementation of OMG’s Model-to-Text standard [11]. MOFScript comes as an Eclipse plug-in. The generated code is then published to the middleware. Figure 14 illustrates the tool chain: The platform-independent adaptation model (PIM) written in UML is transformed via an intermediate transformation into XMI and finally into a platform-specific model (PSM) which in this case targets Java.

Figure 14: Model transformation tool chain

Development of the abstract, platform-independent model and performing automatic code generation provides high flexibility for the development of adaptive applications. If another target-platform should be addressed, the abstract high level model can be reused, only the transformation has to be adjusted to meet the needs of the new platform. Furthermore, if changes in the overall structure of the application are necessary, the modifications can be done at the higher abstraction level of the model and the corresponding code is generated automatically. Abstract system specifications are also very useful to keep track of the set of possible application variants and to ensure completeness.

5 Lessons Learned from Case Studies

Two realistic case studies were developed to evaluate the MADAM approach. Particularly, we wanted to find out about the development overhead for building an adaptive application and the behaviour of the resulting adaptive system as perceived by the end user. The assessment of the development process was based on the comparison of the MADAM platform with commonly used methods, languages and tools. The assessment of the resulting system was based on aspects such as usability, performance, processing overhead and memory footprint.

5.1 Case Studies

In the first scenario, which was derived from a real application of Condat AG written for Daimler Chrysler AG and Deutsche Bahn AG, service technicians are supported during their scheduled tasks regarding inspection, measurements and maintenance of technical equipment such as air conditioning plants or fire alarm installations. The technician uses a PDA for his work. When the technician’s environment changes during his roundtrip, the application on the PDA tries to self-adapt in order to retain its usefulness. For example, in the original version of the application when a communication link deteriorated the technician had to wait until the connection was up again. Using MADAM, the application tries to transmit the data using alternative connections or buffer them temporarily on the local device.

The second pilot application is called SatMotion and was written by Integrasys SA. It is a distributed application which supports the remote line-up of satellite antennas by taking care of power level and receive signal. The adjustment is performed using a PDA that is used to control the measurement equipment and a server over a wireless connection. The application receives the signal traces, visualizes and evaluates them using a spectrum analyzer. The results allow the installer to find the optimal alignment of the satellite antenna. Depending on the installer’s tasks and context conditions, the application assumes different operating modes: the *TwoWay* mode enables to command the remote

instrument and receive signal traces, the *OneWay* mode only receives data for visualization and the Offline mode is able to perform measurements and generate reports on recorded information. If, for example, the network context elements bandwidth and delay deteriorate beyond a certain threshold, the system will switch from *TwoWay* to *OneWay* mode in order to maintain the usability of the application.

In addition, several smaller test applications were written that can be parameterized to measure performance, memory and scalability.

5.2 MDD Approach

The modelling language and tools of the MADAM framework enable the application developer to specify adaptation capabilities at design time and to generate executable code for the runtime environment. The developers model the structure of the pilots in terms of nodes, components, resources, sensors and adaptation capabilities.

Since adaptations in MADAM are defined on the component level, the application has to be divided into a hierarchical collection of components. Alternative functionalities of the application must be located in different components which can be activated or deactivated at runtime as part of the adaptation management. The developer defines a utility function which captures the goals of the adaptation and ensures that the active application configuration always has the highest utility regarding the resource constraints imposed by the environment.

In parallel to the adaptation modelling, the component based application has to be developed, in our case in Java. Obviously, adaptation model and application are highly interrelated, because they refer to the same set of components, context sensors and resources. Especially, the adaptation model defines the composition plans for the exchange of software components in case of context changes. All interrelations must be defined consistently in order to allow a smooth execution of the program during run-time.

The completed model is automatically checked for consistency and transformed into a platform-specific format. The corresponding application is translated into executable code. Both of them are loaded into the MADAM run-time environment for program execution. During run-time the MADAM components access the platform specific model to read the structure and adaptation rules of the loaded application. The adaptation middleware monitors the context and adapts the application in response to context changes.

The development of the adaptive pilot applications required new considerations for the system design. How should the application be divided into components? How can the developer map the intended adaptation strategy onto a utility function? Adding adaptivity to an existing application usually requires a re-engineering of the application architecture. Even if the application is already structured in components, a redesign of the component composition and interfaces is often necessary in order to achieve all needed variants. Since the MADAM approach always implies some overhead in terms of development effort and run-time, it is targeted primarily at distributed applications that have a larger number of context dependencies and adaptation options. One can argue about whether MADAM is beneficial for simple, small applications that depend on only one or two context elements.

During the development of the case studies it turned out that developers found it rather difficult to assign appropriate properties to components and to determine the utility function for an arbitrary application domain. The developer must spend some time to find the right strategy, the property assignments and the utility function for the required adaptations. The best way for the validation and refinement of the selected strategy is often by implementing a pilot application. In addition, the developer needs some experience to find an efficient component architecture, especially for applications with a lot of variants (e.g. > 10.000). A well-chosen component structure and the usage of architectural constraints in the modelling can significantly reduce the amount of variants and improve the performance.

5.3 Prototype Performance

In order to assess the runtime overhead of the MADAM middleware we have performed a set of measurements using the pilot applications and special measurement applications. In this section we

first describe briefly the applied metrics and how the measurements were performed, and then summarise and discuss the main findings.

Relevant metrics were:

- ◆ Memory footprint, i.e. the memory required to store the middleware code and the context, architecture and property prediction models used by the middleware at runtime.
- ◆ Context monitoring overhead, i.e. the additional processing needed to monitor the context and detect context changes.
- ◆ Context latency, i.e. the time from when a context change occurs until it is signalled to the adaptation manager.
- ◆ Planning time, i.e. the time needed by the middleware from when a context change is signalled to the adaptation manager until an action has been decided. During this time the middleware competes with the applications for CPU time.
- ◆ Reconfiguration time, i.e. the time from when the planning has been completed until the new configuration is operational. During this time the applications are blocked. Clearly this time depends on the complexity of the reconfiguration to be done. We measure the time for one component replacement.

In order to ease the isolation of the effect of the middleware from the effect of the normal application execution during measurements, in addition to the two real adaptive applications we developed several special measurement applications that were basically skeletal applications with appropriate models such that they could be adapted by the middleware, but demanded negligible resources to execute. This included

- ◆ an “empty” application, with no variation and no context dependencies,
- ◆ a “tiny” application, with one variation point with two alternatives and one context dependency,
- ◆ a “skeleton” application which is built as a composition of empty components of the same type, which can be scaled in terms of the depth of the composition hierarchy, the number of roles of each composition, and the number of alternative implementations available for each role.

Not all metrics could be measured directly but had to be calculated from more primitive measurable metrics. Also there are sources of error that could disturb the accuracy of the measurements. However, by instrumenting both the pilot applications and the measurement applications appropriately and running a number of test cases repeatedly, we were able to get measures for the metrics we were interested in, which we judge to be sufficiently accurate to assess the feasibility of the MADAM technology. More details about how the measurements were performed can be found in [5].

Figure 15 Testing environment

To aid the measurements we used a scenario engine as illustrated in Figure 15. The scenario engine is driven by scripts and produces simulated context and user input which is fed to the system being tested. The executed scenario and output from the system being tested were recorded in a log file. The hardware environment for the tests consisted of a HP iPAQ PDA and some desktop PCs connected by a wireless LAN. Details are given in Table 1.

Table 1 Devices used for the measurements

<i>Device</i>	<i>Processor</i>	<i>Memory</i>	<i>OS</i>	<i>Java VM</i>
Desktop PC	Pentium 4 2,4 GHz	1 Gb	MS Windows XP SP2	JDK 1.3.1_15
PDA HP iPAQ 6345	Texas Instruments 68 MHz	64 Mb	MS Windows Mobile 2003	CrèMe 4.00

The main results from the performance measurements are shown in Table 2. We performed the measurements on two different configurations, one with the PDA acting as master and two PCs as slaves (column “PDA”), and one with a desktop PC acting as master and two other desktop PCs as slaves (column “Desktop PC”). Obviously, context monitoring overhead and context latency depend on the number and nature of the context properties to be monitored. We measured them for the set of context properties that the pilot applications were sensitive to.

Table 2 Measurement results

<i>Metric</i>	<i>PDA</i>	<i>Desktop PC</i>
Memory (base)	6-8 Mb	6-8 Mb
CPU (continuous)	13%	7%
Context latency (local)	190 ms	23 ms
Context latency (remote)	287 ms	70 ms
Planning (pilot 1, 249 variants)	845 ms	97 ms
Planning (pilot 2, 12 variants)	742 ms	87 ms
Planning (100 variants)	787 ms	94 ms
Planning (1000 variants)	1250 ms	112 ms
Reconfiguration (1 component replacement)	715 ms	94 ms
Reconfiguration (pilot 1, 3 component replacements)	789 ms	112 ms

As already mentioned above, the pilot applications were derived from real products. One has 12 variants and the other one has 249 variants. We believe they are fairly representative for the kind of applications we have to deal with on handheld devices, both in terms of size, context dependencies and adaptation complexity.

The memory footprint clearly depends on the number of active applications and their adaptation complexity. However, the effect was small. The measured memory footprint is in the order of 10% of the available memory on the PDA that was used for the tests, which is not unacceptable,

With one (empty) application running on the handheld device, the extra continuous load on the CPU caused by the MADAM middleware is 13%. The adaptation time (the sum of the planning time and the reconfiguration time), which is the period when the operation of the application may be more severely disturbed, is in the order of a few seconds. We consider this as acceptable, although close to the limit.

On the PDA, one pilot application together with the middleware almost filled up the memory, so scenarios with multiple applications running concurrently are not likely on this class of devices. However, more recent devices offer significantly more resources, typically 10 times more on both processing and memory capacity. On such devices we can envisage scenarios with several applications running concurrently and more complex distributed computing environments. A typical scenario might be three applications similar to our pilot applications, with three auxiliary computers available for running components, the possibility to choose between WLAN, UMTS and GPRS for the communication and three settings for the CPU frequency of the handheld. Extrapolating our measurements to this scenario, we also get adaptation times in the order of a few seconds, which is still acceptable.

In our measurements the planning time and the reconfiguration time contributed approximately on equal terms to the total adaptation time. With a further increase of the number of applications and variation points we quickly move into a steep area of exponential growth of the planning time with the number of variation points, which causes the planning time to dominate and is likely to lead to unacceptable adaptation delays.

5.4 Required Improvements

The validation of the MADAM framework by the pilot applications has shown that the general approach is viable and successful. However, some improvements and enhancements are necessary, which will allow the MADAM middleware and methodology to serve as a foundation for future developments.

The modelling tools and run time environment were evaluated in terms of usability, learning effort, flexibility as well as for the scope and completeness from the developer perspective. One main result of the evaluation was that the modelling approach for self-adaptive mobile systems leads to a systematic and principled way, avoiding partial, ad-hoc solutions. However, the modelling concepts represent sophisticated emerging technology. When designing adaptation models for the pilot applications, the developers asked for more insights into the modelling language and the interrelation between the model elements and the generated code. In a way this contradicts the purity of the model-driven development paradigm but reflects the necessities of real life. In addition, developers wanted more support, i.e. methods and tools, for finding and defining appropriate property predictors and for assigning utility values to component configurations.

Regarding flexibility and completeness, the offered MADAM functionality has proven to be sufficient for implementing distributed and multiple adaptive applications. One potential improvement has been identified: since a fixed load order of distributed MADAM applications is required at present, a more flexible launch should be supported.

Currently, the MADAM middleware prototype allocates about 7 MByte memory, which is acceptable for a PC, but a bit high for a PDA. The performance is good on the PC, but not always on a PDA which runs already loaded large applications. In future work, we intend to improve performance and memory footprint, because this would increase the user acceptance and offers the opportunity to apply MADAM on even smaller devices (e.g. smart phones, multi-functional MP3 players, smart home devices) and to support applications comprising more concurrent processes or variants.

Although the tests indicate that interesting and realistic scenarios can be tackled also on handheld devices without disturbing the normal execution of applications unacceptably, performance improvements both regarding memory footprint and adaptation times are desirable. In particular there is a need to reduce the planning time. An obvious improvement concerns the preciseness of the variability model. In an application normally not all alternatives of all variation points can be combined freely. Typically there are dependencies between variation points, such that choosing one alternative for one variation point constrains the choice at other variation points. This can be expressed as architectural constraints in the MADAM modelling notation but currently the middleware does not exploit this. Therefore “illegal” variants have to be suppressed by ensuring that their utility is zero. This means that during planning the middleware considers more variants than necessary. Often architectural constraints reduce the number of variants to consider substantially. For example, in the pilot with 249 variants the total number of “legal” variants is only 28. Therefore, we believe that leveraging architectural constraints in the middleware may reduce the planning time significantly.

Another possibility is to delegate the planning to a powerful server. Our tests indicate that even on an average laptop, planning is 10 times faster than on the handheld. Delegation of planning will work only when the device is connected, so there must be an onboard backup that takes over planning while the device is disconnected. However, since when a device is disconnected there are normally much fewer application variants that are possible than when it is connected, this approach may also have a significant effect on reducing planning time.

Finally there is the possibility to introduce heuristics in the planning algorithm that is able to make a decision without having to consider all variants. In particular if we relax the requirement to always find the optimal variant, there are opportunities for gaining performance.

In conclusion we think that the development and testing of the pilots have demonstrated the feasibility of the MADAM approach. Adaptations were always meaningful and useful, and the development and runtime overhead were acceptable. In some cases, it would be helpful if significant reconfigurations were indicated suitably to the user, e.g. if an important function is no longer available. Another helpful option would be, if the professional user had means to select variants himself.

6 Related Work

Development methodologies, platforms and middleware supporting dynamic adaptation of context-aware applications on mobile computing devices have been studied extensively during the last decade. In this section we compare the work of MADAM to a few other representative and related efforts. The diversity of research in this area is vast as illustrated in [27] and [2]. For this work we are concerned about related efforts on development methodologies and dynamic middleware platforms for adaptive applications. Clearly, MADAM has not been the first project to tackle dynamically adaptive software systems. It is based on, incorporates, and extends earlier research work. However, MADAM stands out since it provides a truly comprehensive solution that addresses adaptation from both the theoretical and practical perspective and solves many challenges in one single integrated framework.

Early adaptive systems addressed QoS concerns and QoS management involving maintenance and renegotiation of QoS – but these were not really context-aware and not targeting mobile computing [28]. The emergence of adaptive middleware generalized earlier approaches to adaptation by providing open implementations and reflection as a general means to adaptation. Notable examples include the ADAPT project [29] introducing open bindings as a means for adapting video stream bindings, OpenORB [30] introducing meta-level objects and causal connections between the meta and base level in reflective middleware, and dynamicTAO [31] supporting dynamic configuration of the ORB itself. However, these works are still very much QoS oriented focusing on adaptation of middleware level services. In later years, there has been an increasing interest in context-awareness and mobile computing focusing on application-level adaptation. MADAM falls under this category of projects. A comprehensive survey of projects on adaptive software systems is provided in [12]. Below we compare the MADAM framework to some related projects in this domain.

The Bay Area Research Wireless Access Network (BARWAN) project ([40],[41]) at the University of California Berkley, is an early project which combines overlay networks, proxy servers, and reliable data transport for wireless networks. The overlay network provides seamless roaming between mobile communication systems and hand-over between WLAN cells. The proxy servers dynamically adapt retrieved content based on context. In MADAM, an open API was defined for interacting with Mobile IP clients which handle the roaming and handover, and through this API network properties are monitored and collected for use in the adaptation reasoning. MADAM does not predefine the use of proxy servers, but such could be included as part of the architecture of the adaptive applications. In the utility based approach of MADAM, the decision on how to configure a proxy server and whether to use one at all would be made based on what gives the best overall behaviour of the system.

CARISMA [43] is a mobile computing peer-to-peer middleware that exploits the principle of reflection to support the construction of context-aware, adaptive applications. In contrast to MADAM, CARISMA focuses on adaptation of middleware level services. Planning in CARISMA consists of choosing among a handful of predefined rule based adaptation policies (referred to as application profiles) using utility functions and resolving conflicts between policies using an auction-like procedure. While the adopted rule based policies are natural and simple to use, there are some drawbacks compared to MADAM's use of utility functions. Firstly, rules provide less transparency to

the developer by requiring reasoning in terms of lower-level reconfiguration actions rather than architectural design with variation points. In MADAM, the lower level reconfiguration actions are automatically determined by the middleware. Adaption in MADAM is QoS (or property) driven while the CARISMA rule based policies do not consider prediction of non-functional properties. Lastly, CARISMA handles only a fixed number of adaptation policies whereas the utility function approach of MADAM allows choosing the best adaptation among an open-ended range of alternatives, both fine grained and course grained, that are automatically derived by the adaptation middleware itself.

Chisel [44] is an open framework for dynamic adaptation of services using reflection in a policy-driven, context-aware manner. It is based on decomposing the extra-functional aspects of a service into alternative behaviours. Then, a policy driven algorithm is used to dynamically select the most suitable alternative, based on the context changes. Similar to MADAM, context changes trigger a process that intelligently determines if an adaptation is required. Unlike MADAM which uses utility functions, Chisel leverages human-readable declarative adaptation policies to instruct the adaptive behaviour of the applications. These policies are defined with a custom-made language and in contrast to MADAM are essentially rule based

MobiPADS [45] enables active service deployment and reconfiguration in response to a dynamic environment. This is achieved with the use of so-called mobilets, which are active pieces of mobile code existing in pairs: a primary mobilet executing on the client (mobile) node, and a secondary mobilet executing on the server side. The mobile client offloads as much of the computation and storage burden as possible to the server side. Once a context change is detected, the MobiPADS middleware either reconfigures the current mobilet service chain, by notifying the mobilets about the context changes so they can internally adapt themselves, or by instructing the applications to adapt. While the intentions of MobiPADS are similar to the ones of MADAM, obviously they have chosen a different technical approach.

Odyssey [42] is a QoS-middleware that dynamically chooses a version of the content before transporting it across the wireless network. Odyssey is running on the client side and the content is stored on the server side. Applications on the mobile computer run independently from each other. The adaptation decisions depend primarily on the quality of the transmission link. Applications provide a window of tolerance, i.e. a lower and upper bound on the required resource quantities. If resource capacity changes beyond these thresholds, the application will be notified. Application specific Wardens capture the adaptation possibilities of adaptable applications and offer a choice of adaptation levels. Rather close to the principle of the utility function of MADAM, Odyssey introduces the idea of fidelity to label different versions of the same content stored on a server. Fidelity is the degree to which the content presented matches the original reference copy. A fidelity level is selected by mapping monitored QoS values of the network to fidelity levels.

Quality Objects (QuO) [36, 37] is a well-known middleware project. Although QuO does not address the fundamental problems associated with portable devices and wireless networks, it has had a major impact on adaptive and QoS-aware middleware development projects, including MADAM. QuO is an architecture that focuses on providing design-time and run-time support for adaptive client-server applications. The QuO approach enables applications to adapt their functionality dynamically to changing communication and computing resources. It requires run-time monitoring of the resources and a detailed description of all possible adaptations in each application component. QuO contracts define the possible application adaptations under changing resource characteristics. QuO contracts define a QoS space that is divided into regions that are defined by predicates on system properties. When system properties change, QoS can move from one region to another, which may trigger adaptive actions to keep a certain quality level. A disadvantage of the QuO architecture is that it only focuses on adapting the application's internal functionality based on the availability of system resources. QuO does not consider adaptations in the deployment configuration of a distributed application. QuO (in addition to MADAM) is also one of a few projects that have attempted to address adaptation as a reusable component or service. Specialized QoS management functions are encapsulated by so-called Qoskets. Qosket components provide a general solution to manage end-to-end QoS requirements without requiring the middleware to provide such services. However, QuO does not provide a framework that is able to reason about the needs for QoS management functions at run-time, and to map these needs to concrete Qosket components, and patterns for how to use them.

The QuA project [38] has developed a QoS-aware adaptation middleware that is based on many of the same principles as the MADAM adaptation management, i.e., using run time models, utility functions, QoS prediction and service planning. In contrast to MADAM, QuA provides a generic middleware level service that can be used to retrieve reflective models, referred to as service mirrors of applications and services in any phase of their life cycle, including run-time. Service mirrors are maintained by the middleware to reflect the life-cycle state of the entities they model. QuA does not provide a comprehensive context management middleware but may exploit a context middleware similar to the one provided by MADAM. The QuA middleware does not focus on mobile applications, but rather on supporting adaptive multimedia streaming applications. A related effort, QuAMobile [39] has tuned the QuA core middleware towards context-aware mobile computing systems. QuAMobile introduces a QoS-architecture and a distributed meta-level representation of context and services for mobile computing. The QoS architecture of QuAMobile has been validated as a proof-of-concept implementation, and thus no complete prototype has till date been realized.

The self-adaptation techniques proposed by the Rainbow approach [20] are similar to MADAM in that they attempt to separate adaptation from application logic concerns. Rainbow extends architecture-based adaptation by adopting architectural styles. Architectural styles are used to tailor the adaptation infrastructure by encoding the developer's system-specific knowledge, including adaptation strategies and system concerns. However, Rainbow makes no attempt to tackle mobile environment requirements. Furthermore, its developers appear to have based the adaptation strategies on situation-action rules, which specify exactly what to do in certain situations. In contrast, MADAM uses extended goal policies expressed as utility functions, which is a higher level of adaptation strategy specification that establishes behavioural objectives, and which leaves the system to reason on the actions required to implement those policies.

ReMMoC [34] is a dynamic middleware that supports interoperability between mobile clients and ubiquitous services. During run-time, the ReMMoC service discovery component reconfigures itself and the remote service binding to match the protocols of the discovered ubiquitous services. ReMMoC provides an abstract programming model that alleviates some of the problems associated with designing applications for dynamic middleware platforms. The programming model offers generic APIs that are mapped to the technology specific APIs. Like MADAM, ReMMoC uses architecture specifications for both the initial configuration and reconfigurations. ReMMoC does not support anything like service planning or search for service implementation alternatives, but applies rule-based policies limited to a fixed set of static component compositions.

UbiQoS [35] is a context-aware middleware that adapts media content according to QoS requirements and portable device characteristics. The middleware assumes a proxy and a QoS-adapter in the wireless network. QoS requirements of different users and context information (including network QoS) about the mobile terminals are stored in a centralized server. This information is used by the adapter to adapt the content to fit the portable device and the network QoS, and to satisfy the user QoS requirements. UbiQoS does not implement anything like generalized run time service planning as in MADAM but focuses on QoS maintenance of media content. Furthermore, UbiQoS implements context sensing, but does not provide context information management interfaces to applications or other middleware functions as in MADAM.

Regarding the development methodology for context-aware, adaptive applications, the research reported in [32] and [33] has some similarities to the MADAM approach. Their approach also emphasizes the importance of models. A rich set of conceptual models for modelling context dependencies and adaptation is proposed to support the software engineering process. Adaptation is controlled by defining context-dependent preferences for the application behaviour. A software infrastructure is described that supports the context-awareness and adaptability of applications. The overall objectives are almost identical to MADAM. However, MADAM has chosen a development methodology that builds heavily on the notion of dynamically reconfigurable component frameworks, adaptation decisions based on dynamically computed utility functions, and an MDA-compliant development process based on UML models and popular MDA tools. Furthermore, MADAM supports unanticipated adaptations by enabling the introduction of new components at run-time.

7 Conclusions and Future Work

Run-time adaptation will be a major requirement for future software systems in ubiquitous computing environments. Technologies such as mobile computing, ad-hoc networking, pervasive computing infrastructures, and dynamic service discovery clearly contribute to this trend. The MADAM project has shown the viability of a general and comprehensive approach to the development of self-adaptive applications for mobile computing scenarios. MADAM provides a computing infrastructure as well as a development methodology for novel applications that are able to react to context changes and adapt their behaviour in order to maintain their usefulness under different execution conditions. Self-adaptation will not only lead to more useful and more versatile, but also to more robust and more dependable applications.

On the one hand MADAM has tackled successfully a rather broad range of research issues. On the other hand it has opened up our eyes for a number of new challenges that wait to be solved. For example, usability and human-computer interaction are important issues where we need more insights. Will the user accept systems that adapt automatically? How often may an application adapt without becoming an annoyance to the user? Does the user trust a system that changes automatically? This leads to another open problem area: How do you validate and test a self-adaptive software system that is able to handle unanticipated adaptation at run-time? How about security and trust in distributed context information? Many of these questions can only be answered through experiments and practical experience with self-adaptive systems.

Furthermore, since the inception of MADAM several years ago new technological achievements have introduced additional requirements and opportunities. For example, ad-hoc networking facilities and ubiquitous service architectures are made available that represent an enrichment of an application's execution context. Thus, an adaptive application may want to replace a local component by a remote service if it promises a better service. We need new context models and context query languages to model these environments and fully exploit such scenarios. The usage of remote services must be controlled by some implicit or explicit service level agreement. Adaptation decisions may depend on the quality of a service as well as on its price. It has to be explored what kind of decision support techniques are appropriate for controlling such adaptations. Even more open questions arise when we also allow the adaptation middleware to adapt dynamically.

We have recently started a new research project called MUSIC which addresses many of these open questions by building on the results of MADAM and extending them towards ubiquitous computing scenarios. Thus, the ideas underlying the MADAM project represent a long-term and continuing research thread in the realm of self-adaptive software.

ACKNOWLEDGEMENTS

The contributions of all members of the MADAM team are gratefully acknowledged. We thank the EC for the financial support under contract FP6 IST 4159.

REFERENCES

1. MADAM Project Home Page, <http://www.intermedia.uio.no/confluence/display/madam>
2. Software: Practice and Experience, Special Issue: Experiences with Auto-adaptive and Reconfigurable Systems, Edited by M. Aksit, T. Elrad Volume 36, Issue 11-12, John Wiley & Sons (2006)
3. MADAM Deliverable D1.2, Adaptation Scenarios and Reference Requirements for Service Adaptation, A. Chimaris, G. Papadopoulos, N. Paspallis, Z. Abraham (editors), <http://www.intermedia.uio.no/confluence/display/madam/Deliverables>
4. MADAM Deliverable D5.3, Pilot Services Adaptation Requirements and Design, P. Ruiz, J. Sanchez, R. Fricke (editors), <http://www.intermedia.uio.no/confluence/display/madam/Deliverables>
5. MADAM Deliverable D6.2, Pilot Services Trial Evaluation, Rolf Fricke, John Klimasek, Ali Sarioglu (editors), <http://www.intermedia.uio.no/confluence/display/madam/Deliverables>

6. M. Alia, G. Horn, F. Eliassen, M.U. Khan, R. Fricke and R. Reichle, "A Component-based Planning Framework for Adaptive Systems". The 8th International Symposium on Distributed Objects and Applications (DOA), Oct 30 - Nov 1, 2006, Montpellier, France.
7. Mohammad Ullah Khan,, Roland Reichle, Kurt Geihs, Applying Architectural Constraints in the Modeling of Self-adaptive Component-based Applications, Workshop on Model Driven Software Adaptation (M-ADAPT'07) at 21st European Conference on Object Oriented Programming (ECOOP), 2007, Berlin, Germany
8. Kurt Geihs, Mohammad U. Khan, Roland Reichle, Arnor Solberg, Svein Hallsteinsen, "Modeling of Component-Based Self-Adapting Context-Aware Applications for Mobile Devices". IFIP Working Conference on Software Engineering Techniques (SET), 2006, Warsaw, Poland.
9. MOFScript Model-to-Text Transformation, <http://www.eclipse.org/gmt/mofscript/>
10. XML Metadata Interchange (XMI), v2.1, <http://www.omg.org/cgi-bin/doc?formal/2005-09-01>
11. MOF Models to Text Transformation Language, Beta 2, OMG Document Number: ptc/07-08-16, <http://www.omg.org/docs/ptc/07-08-16.pdf>
12. P. K. McKinley, S. M. Sadjadi, E. P. Kasten, and B. H. C. Cheng, Composing Adaptive Software, IEEE Computer, 37(7):56-64, July 2004.
13. MADAM Deliverable D3.3, UML Modelling Elements and Approach for Application Adaptation, Kurt Geihs, Mohammad Ullah Khan, Roland Reichle (editors), <http://www.intermedia.uio.no/confluence/display/madam/Deliverables>
14. Anind K. Dey, "Understanding and Using Context", Personal and Ubiquitous Computing, Vol. 5, No. 1, 2001, pp. 4-7
15. Guanling Chen, David Kotz, "A Survey of Context-Aware Mobile Computing Research", Technical Report TR2000-381, Department of Computer Science, Dartmouth College, Nov. 2000
16. M. Mikalsen, N. Paspallis, J. Floch, G.A. Papadopoulos, P.A. Ruiz, "Putting Context in Context: The Role and Design of Context Management in a Mobility and Adaptation Enabling Middleware", Intern. Workshop on Managing Context Information and Semantics in Mobile Environments (MCISME) in conjunction with the 7th International Conference on Mobile Data Management (MDM), Nara, Japan, May 9-12, 2006, IEEE Computer Society Press, pp.76-83
17. Nearchos Paspallis, Avraam Chimaris, George A. Papadopoulos, Experiences from Developing a Context Management System for an Adaptation-enabling Middleware, 7th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS), Paphos, Cyprus, June 5-8, 2007, Springer LNCS 4531, pp. 225-238
18. Floch, J., E. Stav, and S. Hallsteinsen. Interfering effects of adaptation: implications on self-adapting systems architecture. In proceedings of Distributed Applications and Interoperable Systems: 6th IFIP WG 6.1 International Conference, DAIS 2006, Bologna, Italy, June 14-16, 2006. Springer-Verlag GmbH. p. 64-69.
19. Oreizy, P., et al. Architecture-based approach to self-adaptive software. IEEE Intelligent Systems and Their Applications, 1999. 14(3): p. 54-62.
20. Garlan, D., et al. Rainbow: Architecture-based self-adaptation with reusable infrastructure. Computer, 2004. 37(10): p. 46-54.
21. Lutfiyya, H., et al. Issues in Managing Soft QoS Requirements in Distributed Systems Using a Policy-Based Framework. In proceedings of Proc. 2nd International Workshop on Policies for Distributed Systems and Networks (POLICY '01). 2001. Springer-Verlag. p. 185--201.
22. Kephart, J.O. and D.M. Chess. The vision of autonomic computing. IEEE Computer, 2003. 36(1): p. 41-52.
23. W.E. Walsh et al., "Utility Functions in Autonomic Systems," Proc. Int'l Conf. Autonomic Computing, IEEE CS Press, 2004, pp. 70-77.
24. Kephart, J.O., and Das, R. Achieving Self-Management via Utility Functions. IEEE Internet Computing, vol.11, no. 1 , pp. 40-48, 2007.
25. Poladian, V., et al. Dynamic configuration of resource-aware services. In Proc. of 26th Intern. Conf. on Software Engineering. 2004. Edinburgh, UK. IEEE Computer Society p. xviii+786.
26. Floch, J., et al. Using architecture models for runtime adaptability. IEEE Software - Special Issue on Software Architecture, 2006. 23(2): p. 62-70.
27. Mascolo, C., Capra, L., Emmerich, W., "Mobile Computing Middleware", Advanced Lectures in Network, Lecture Notes in Computer Science, Vol.2497, pp.20-58]

28. Cristina Aurrecochea, Andrew T. Campbell, Linda Hauw, A survey of QoS architectures Multimedia systems, 6: 138-151. 1998].
29. Tom Fitzpatrick, Gordon Blair, Geoff Coulson, Nigel Davies, Philippe Robin, "Supporting Adaptive Multimedia Applications through Open Bindings",. International Conference on Configurable Distributed Systems (ICCDs '98) Annapolis, Maryland, USA May 1998]
30. Blair G., Coulson G., Robin P., Papathomas M., "An Architecture for Next Generation Middleware" , in Middleware '98 IFIP Intl. Conf. on Distributed Systems Platforms and Open Distributed Processing, Springer 1998, pp. 191-206]
31. F. Kon, M. Roman, P. Liu, J. Mao, T. Yamane, L.C. Magalhães, R.H. Campbell, "Monitoring, Security, and Dynamic Configuration with the dynamicTAO Reflective ORB", Middleware'2000, New York, USA, April 2000.]
32. K. Henriksen and J. Indulska, "Developing context-aware pervasive computing applications: Models and approach", Pervasive and Mobile Computing, Vol. 2, No. 1 , February 2006, p. 37-64
33. K. Henriksen, J. Indulska, A. Rakotonirainy, "Using context and preferences to implement self-adapting pervasive computing applications", Software Practice and Experience, Vol. 36, No. 11-12, Feb. 2006, p. 1307-1330
34. Grace, P., Blair, G., Samuel, S.: ReMMoC: A Reflective Middleware to Support Mobile Client Interoperability, *Proc. of Intern. Symposium on Distributed Objects and Applications*, Lecture Notes in Computer Science, Vol. 2888, Springer-Verlag, 2003, pp. 1170-1187
35. P. Bellavista, C. Stefanelli, M. Tortonese, "The Ubiquitous provisioning of Internet services to portable devices, IEEE Pervasive Computing, 1(3), 2002, pp. 81-87.
36. Heineman, G. T., J. P. Loyall, et al. (2004), Component Technology and QoS Management, Intern. Symposium on Component-based Software Engineering (CBSE7), Edinburgh, Scotland.
37. Sharma, P.K., Loyall, J.P., Heineman, G.T., Schantz, R.E., Shapiro, R., and Duzan, G., "Component-Based Dynamic QoS Adaptations in Distributed Real-Time and Embedded Systems". In *Proceedings of the International Symposium on Distributed Objects and Applications (DOA'04)*, 2004.
38. Gjorven, E., Eliassen, F., Lund, K., Eide, V. S. W., and Staehli, R., "Self-Adaptive Systems: A Middleware Managed Approach", In *Proceedings of the 2nd IEEE International Workshop on Self-Managed Networks, Systems and Services (SelfMan'06)*, 2006.
39. Lundesgaard, S.A., Lund, K, Eliassen, F.: Utilising Alternative Application Configurations in Context- and QoS-aware mobile middleware, *Proc. of the 6th IFIP WG 6.1 Intern. Conference on Distributed Applications and Interoperable Systems*, Springer-Verlag, 2006, pp. 228-241
40. Katz, R.H., Brewer, E.A., Amir, E.; Balakrishnan, H., Fox, A.; Gribble, S., Hodes, T., Jiang, D., Giao Thanh Nguyen, Padmanabhan, V., Stemm, M. "The Bay Area Research Wireless Access Network (BARWAN)". COMPCON Spring '96 - 41st IEEE Intern. Computer Conference, 1996.
41. Brewer, E.A., Katz, R.H., Chawathe, Y., Gribble, S.D., Hodes, T., Giao Nguyen, Stemm, M., Henderson, T., Amir, E., Balakrishnan, H., Fox, A., Padmanabhan, V.N., Seshan, S. "A network architecture for heterogeneous mobile computing". Personal Communications, IEEE, vol.5, no.5, pp.8-24, Oct 1998.
42. Noble, B. D. and M. Satyanarayanan. "Experience with adaptive mobile applications in Odyssey". Mobile Networks and Applications 4(4): 245-254, 1999.
43. Licia Capra, Wolfgang Emmerich, and Cecilia Mascolo, "CARISMA: Context-Aware Reflective Middleware System for Mobile Applications", IEEE Transactions on Software Engineering, Volume 29, Issue 10, October 2003, pp. 929-945
44. John Keeney, Vinny Cahill, "Chisel: A Policy-driven, Context-aware, Dynamic Adaptation Framework", 4th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2003), Lake Como, Italy, June 4-6, 2003, pp. 3-14
45. Alvin T.S. Chan, Siu-Nam Chuang, "MobiPADS: A Reflective Middleware for Context-Aware Mobile Computing", IEEE Transactions on Software Engineering (TSE), Volume 29, Issue 12, Dec. 2003, pp. 1072 – 1085