

Kent Academic Repository

Full text document (pdf)

Citation for published version

Efstratiou, Christos and Cheverst, Keith and Davies, Nigel and Friday, Adrian (2001) An Architecture for the Effective Support of Adaptive Context-Aware Applications. In: MDM '01. Proceedings of Mobile Data Management (MDM'01). Springer, Berlin pp. 15-26. ISBN 3-540-41454-1.

DOI

Link to record in KAR

<https://kar.kent.ac.uk/38869/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

An Architecture for the Effective Support of Adaptive Context-Aware Applications

Christos Efstratiou, Keith Cheverst, Nigel Davies and Adrian Friday

Distributed Multimedia Research Group, Department of Computing, Lancaster University,
Bailrigg, Lancaster LA1 4YR, U.K.
e-mail: most@comp.lancs.ac.uk

Abstract. Mobile applications are required to operate in environments characterised by change. More specifically, the availability of resources and services may change significantly during typical periods of system operation. As a consequence, adaptive mobile applications need to be capable of adapting to these changes to ensure they offer the best possible level of service to the user. Our experiences of developing and evaluating adaptive context-aware applications in mobile environments has led us to believe that existing architectures fail to provide the necessary support for such applications. In this paper, we discuss the shortcomings of existing approaches and present work on our own architecture that has been designed to meet the key requirements of context-aware adaptive applications.

1 Introduction

Mobile applications are required to operate in environments that change. Specifically, the availability of resources and services may change significantly and frequently during typical system operation [8, 11]. As a consequence, mobile applications need to be capable of adapting to these changes to ensure they offer the best possible level of service to the user [11]. While early research focused on applications which adapted to changes in network characteristics, there is now increasing interest in applications that adapt to general environmental and contextual triggers such as changes in a system's physical location, e.g. the GUIDE system [2, 3] which supplies users with information tailored to their current location.

Current adaptive mobile applications are built using one of two approaches: either the adaptation is performed by the system which underpins the application (in an attempt to make transparent the effects of mobility) or the application itself monitors and adapts to change. In some cases, these approaches are combined. For example, in the MOST system [8] where the middleware platform adapts the operation of the network protocol in the face of changes in QoS and, additionally, reports these changes to the application to enable application-level adaptation. However, in the general case, it has been demonstrated that maintaining transparency in the face of mobility is not practical and that it is difficult for a system to adapt without support from the application.

Specifically, the availability of resources and services may change significantly and frequently during typical system operation. As a consequence, mobile applications need to be capable of adapting to these changes to ensure they offer the best possible level of service to the user. Our experiences of developing and evaluating adaptive context-aware applications in mobile environments has led us to believe that existing architectures fail to provide the necessary support for such applications. In this paper, we discuss the shortcomings of existing approaches and present work on our own architecture that has been designed to meet the key requirements of context-aware adaptive applications.

Specifically, the availability of resources and services may change significantly and frequently during typical system operation. As a consequence, mobile applications need to be capable of adapting to these changes to ensure they offer the best possible level of service to the user. Our experiences of developing and evaluating adaptive context-aware applications in mobile environments has led us to believe that existing architectures fail to provide the necessary support for such applications. In this paper, we discuss the shortcomings of existing approaches and present work on our own architecture that has been designed to meet the key requirements of context-aware adaptive applications.

Careful examination of current approaches to supporting adaptation reveal two important facts. Firstly, support for adaptation is often fragmented with a range of mechanisms being used to notify applications of changes in different environmental and contextual attributes [4]. Secondly, there is a lack of mechanisms that support coordination of adaptive behaviour across the whole system, according to user requirements. In this paper, we explore the requirements for, and our research into creating a unified architecture to support multiple contextual attributes coupled with a user-driven adaptation control mechanism. The benefits of such an approach are clearly illustrated using set of real-world examples.

2 Drawbacks of Current Approaches

Mobile systems need to be capable of adapting to a wide range of attributes such as network bandwidth, location, power etc. In general, current mobile systems provide support for adaptive applications by notifying applications when certain 'interesting' changes in attributes occur, e.g. bandwidth falls below some specified minimum threshold. It is then the responsibility of the application to adapt in an appropriate way, e.g. by reducing its bandwidth requirements. However, this approach can be shown to lead to inefficient solutions because of a lack of support for enabling multiple applications that may co-exist on the same system. In the following scenario, we illustrate the kind of problems that could occur as a result of relying on a simplistic notification based approach and isolated uncoordinated application adaptation.

Mobile systems need to be capable of adapting to a wide range of attributes such as network bandwidth, location, power etc. In general, current mobile systems provide support for adaptive applications by notifying applications when certain 'interesting' changes in attributes occur, e.g. bandwidth falls below some specified minimum threshold. It is then the responsibility of the application to adapt in an appropriate way, e.g. by reducing its bandwidth requirements. However, this approach can be shown to lead to inefficient solutions because of a lack of support for enabling multiple applications that may co-exist on the same system. In the following scenario, we illustrate the kind of problems that could occur as a result of relying on a simplistic notification based approach and isolated uncoordinated application adaptation.

2.1 Scenarios

The Need for Coordinated Application Adaptation for Power Management. This scenario illustrates the need for coordination in order to achieve efficient power management on a mobile system. One existing approach for handling power management, i.e. the ACPI [1] model, is to enable the operating system to switch hardware resources into low power mode when not in use, e.g. spinning down the hard-disk. This approach requires that applications leave hardware resources in an idle state for sufficient period of time to make a transition between idle and active states worthwhile. Although this approach is suitable when only one application is running on a mobile device, the approach can prove ineffective when multiple applications or system services are sharing hardware resources. In more detail, the lack of coordinated access to hardware resources can result in poor utilisation of the shared resource and therefore sub-optimum power management. For example, consider the case of multiple applications that implement a auto-save feature. If the absence of any coordination between applications results in each application saving its state to the disk at an arbitrary time, then the application may choose to save its state to the disk without considering the state of other applications. If the applications are able to coordinate their access to the hard-disk then access to the disk can be clustered, allowing longer periods of inactivity. This latter approach is clearly more power efficient than the situation in which usage of the hard-disk is completely arbitrary and uncoordinated.

The Need for Coordinated Application Adaptation for Power Management. This scenario illustrates the need for coordination in order to achieve efficient power management on a mobile system. One existing approach for handling power management, i.e. the ACPI [1] model, is to enable the operating system to switch hardware resources into low power mode when not in use, e.g. spinning down the hard-disk. This approach requires that applications leave hardware resources in an idle state for sufficient period of time to make a transition between idle and active states worthwhile. Although this approach is suitable when only one application is running on a mobile device, the approach can prove ineffective when multiple applications or system services are sharing hardware resources. In more detail, the lack of coordinated access to hardware resources can result in poor utilisation of the shared resource and therefore sub-optimum power management. For example, consider the case of multiple applications that implement a auto-save feature. If the absence of any coordination between applications results in each application saving its state to the disk at an arbitrary time, then the application may choose to save its state to the disk without considering the state of other applications. If the applications are able to coordinate their access to the hard-disk then access to the disk can be clustered, allowing longer periods of inactivity. This latter approach is clearly more power efficient than the situation in which usage of the hard-disk is completely arbitrary and uncoordinated.

The Problem of Conflicting Adaptation.

In this scenario, we illustrate the potential problems that can occur in a system that utilises different attributes. We consider a hypothetical mobile system which utilises two independent adaptation mechanisms, one for managing power and the other for managing network bandwidth. These two mechanisms can conflict with one another as the following example illustrates. If the system needs to reduce power consumption, the power management mechanism will request those network devices to postpone their use of the network bandwidth to enter sleep mode. As a consequence, the use of the network bandwidth adaptation mechanism will detect this unused bandwidth and utilise the spare bandwidth. In this way, the request to postpone network

directly conflict with the request to postpone network adaptation mechanisms. Clearly, coordination is required to detect and avoid potentially conflicting adaptation mechanisms. In the example presented, the instruction given to applications to be withheld from serving power was the system's

In this scenario, we illustrate the potential separate adaptation mechanisms for a mobile system which utilises two independent adaptation mechanisms, one for managing power and the other for managing network bandwidth. These two mechanisms can conflict with one another as the following example illustrates. If the system needs to reduce power consumption, the power management mechanism will request those network devices to postpone their use of the network bandwidth to enter sleep mode. As a consequence, the use of the network bandwidth adaptation mechanism will detect this unused bandwidth and utilise the spare bandwidth. In this way, the request to postpone network

independent and uncoordinated adaptation mechanisms is required in order to detect and avoid potentially conflicting adaptation mechanisms. In the example presented, the instruction given to applications to be withheld from serving power was the system's

Utilising Alternative Location Sensing Mechanisms.

This scenario considers the case of supporting multiple services for providing this case, we consider a location aware system that location through two different mechanisms: local in-cell-based wireless network. Using the latter the current cell in which it operates and thus special deal with the same problem but they both have different characteristics. The GPS mechanism is typically more accurate (accuracy requires extra power to operate. Alternatively, the less accurate (depending on the size of a cell, e.g. WaveLAN). However, the fact that the network device for communications means that the addition power consumption required for identifying the location of the base-station would

be minimal. If it follows that the mobile system would select the more accurate location information was required and once was not an important issue. Alternatively, if the time of the system's battery (and therefore operation) was more important than achieving greater location accuracy, the network location mechanism would be the most appropriate adaptive strategy that would be most appropriate depending on the user's requirements and the context of other attributes, the system to make such decisions there is a basic requirement for system-wide adaptation policies. Without such policies, coordinating alternative context retrieval mechanisms is difficult as a single mechanism without being able to identify other system resources and consequently the app

This scenario considers the similar contextual information. In a capable of sensing its current GPS device and using the beaconing mechanism, the system can identify its location. Both mechanisms have different characteristics. The GPS in the region of 5m) but does not require extra power to operate. Alternatively, the less accurate (approximately 200m for WaveLAN). However, the fact that the network device for communications means that the addition power consumption required for identifying the location of the base-station would

be minimal. If it follows that the mobile system would select the more accurate location information was required and once was not an important issue. Alternatively, if the time of the system's battery (and therefore operation) was more important than achieving greater location accuracy, the network location mechanism would be the most appropriate adaptive strategy that would be most appropriate depending on the user's requirements and the context of other attributes, the system to make such decisions there is a basic requirement for system-wide adaptation policies. Without such policies, coordinating alternative context retrieval mechanisms is difficult as a single mechanism without being able to identify other system resources and consequently the app

3 Analysis

Multiple Attributes. The previous scenarios illustrated a number of potential problems with current approaches to developing adaptive context-aware mobile systems. In this section, we generalise on these findings to present a critique of existing mobile systems and their suitability for supporting adaptive applications.

Based on the ideas of ubiquitous computing [20] it is possible to discover changes in both the user and system environment and adapt these changes. Current context-aware applications handle context in an improvised fashion. Application developers usually bundle the application with specific mechanisms for accessing context. However, this approach does not allow coordinated adaptation on context changes leading to the problems presented. Dey [5] has addressed this problem and suggested a general platform to support context-aware applications. Our belief is that though a general platform for supporting context-aware applications is necessary, this platform should also be capable of addressing the problems of coordinated adaptation.

This situation is complicated still further by the fact that the adaptive behaviour of the attributes. These side-effects could, in-turn, trigger adaptation requests to other applications that result in conflicting actions (as illustrated in the conflict in adaptation scenarios in section 2.1.2). Moreover, current research [4, 6, 7, 12, 13] has identified the need to provide adaptation solutions based on the combination of different attributes.

Existing architectures do not provide the necessary support to enable programmers to construct applications which adapt to multiple attributes and identify and cope with conflicting adaptation strategies.

Adaptation Mechanism. Current mobile systems supporting adaptive applications tend to rely heavily on integrating QoS feedback and adaptation with network bindings. Examining the architecture of such systems allows us to identify a framework for analysing the architectural model of existing adaptive systems. The framework comprises two layers, the upper application layer and the lower representing the adaptation support platform. Between these two layers we can identify four distinct flows of control and information (see figure 1).

Flow A. Represents the requirements set by the application concerning the resources or attributes supported by the underlying infrastructure. For example, in the case of network adaptation this flow could represent the application's network QoS requirements.

Flow B. Represents the ability of the application to control the functionality of the underlying infrastructure. In the case of accessing a GPS device this could represent, for example, the control of the device by the application.

Flow C. Represents an information flow from the platform to the application. This could be used, for example, as a notification mechanism to inform the application when certain requirements cannot be met. Such notification could then trigger the application to adapt.




Fig1. Directed flow between applications and platform

FlowD. Represents the ability of the underlying platform to actually control the operation of the application. More specifically, this flow represents an explicit request from the system for the application to perform a specific adaptive behaviour. For example, the application might be requested to reduce its demand for network bandwidth or disk usage.

Consideration of this framework enables a classification of current systems according to the types of flows supported. For example, network-based adaptive systems such as BAYOU [18], Odyssey [16], MOST [8] and Rover [10] support flows A and C.

Context-aware applications like GUIDE [2,3], Stick-Note [14] and Cyberguide [17] are based on flows B and C. In more detail, flow B represents the access to the various context-sensor information flowing from the sensors to the application, while flow C represents the

information flowing from the sensors to the application. According to our knowledge, no platforms supporting context-aware adaptation provides a flow of control from the platform to the application. Indeed, although examples of systems providing this type of flow can be found in the distributed systems community, e.g. ISIS-META [15], it should be noted that these systems consider only network-triggered adaptation.

4 Architectural Requirements

The previous sections have described the limitations of current approaches for supporting adaptive mobile applications. In particular, these approaches do not provide appropriate support to enable applications to adapt multiple attributes in an efficient and coordinated way. This section considers requirements that could be used to develop an appropriate architecture for supporting adaptive mobile applications.

4.1 Supporting Common Space for Extensible Attributes

The first key requirement of the architecture is to provide a common space for handling the adaptation attributes used by the system. It is important that new attributes are not introduced into the system and when they become important, e.g. information about human physiology for wearable computers. The fact that new contextual attributes for triggering adaptation can arise implies that:

- these attributes that trigger adaptation need to be extensible,
- the characteristics of all these attributes may vary.

4.2 Application Control and Coordination

A secondary architectural requirement is the need to support the control of adaptive behaviour across all components involved in that of the main limitations of current approaches is that responsible for triggering an adaptive mechanism which notifies them about any changes. In order to support adaptation there is a requirement for the trigger level. Given this approach, the decision about when to adapt is pushed into an external entity, with cross-adaptive behaviour still part of the application.

As described earlier, one application itself is not the underlying infrastructure that flexible and coordinated adaptation on a system-wide and how applications should be implemented, while the application knowledge, while the application characteristics.

4.3 Support for System Wide Adaptation Policies

A further requirement is to support the notion of system-wide adaptation policies. More specifically, such policies should enable more than one system to operate differently based on the current context and requirements of the user.

system-wide adaptation policies. More specifically, such policies should enable more than one system to operate differently based on the current context and requirements of the user.

The specification of adaptation policies should be goal-oriented. Two kinds of policies can be identified:

goal-oriented. Two kinds of policies can be identified:

1. effects on resources. The policy specifies a specific aim for the use of a specific resource. Example policies include reducing the requirement for network bandwidth and maximising the duration of operation of the system.
2. effects on applications. The policy specifies a mode of operation for specific applications. Example policies include defining priorities on applications which determine the order in which they are allocated resources and maximising the duration of operation of the system while having a specific application operating with full functionality.

1. effects on resources. The policy specifies a specific aim for the use of a specific resource. Example policies include reducing the requirement for network bandwidth and maximising the duration of operation of the system.
2. effects on applications. The policy specifies a mode of operation for specific applications. Example policies include defining priorities on applications which determine the order in which they are allocated resources and maximising the duration of operation of the system while having a specific application operating with full functionality.

5. Architecture

5.1 Structure of the Platform

We propose that future mobile adaptive applications which mechanisms and policies are decoupled and fully exposed and externalised in order to enable control architectures that have been designed to address these requirements.

should adopt an architecture in which the mechanisms can be by independent entities. Our requirements.

Figure 2 shows the relationship between the main components of our proposed platform. The basic functionality of the architecture is two-fold, namely the discovery and control of services offering contextual information and the coordination of adaptive behaviour of the system based on changes in context. More specifically, the platform discovers available context information in the system's environment and manipulates the contextual information in the *context database*.

components of the architecture. The architecture is two-fold, namely the discovery and control of services offering contextual information and the coordination of adaptive behaviour of the system based on changes in context. More specifically, the platform discovers available context information in the system's environment and manipulates the contextual information in the *context database*.

Context aware applications expose their adaptive mechanism to the platform by registering with the *application database*. The *adaptation control* driven by *adaptation policies* (as specified by the user) coordinates the coexistence of applications according to changes in context.

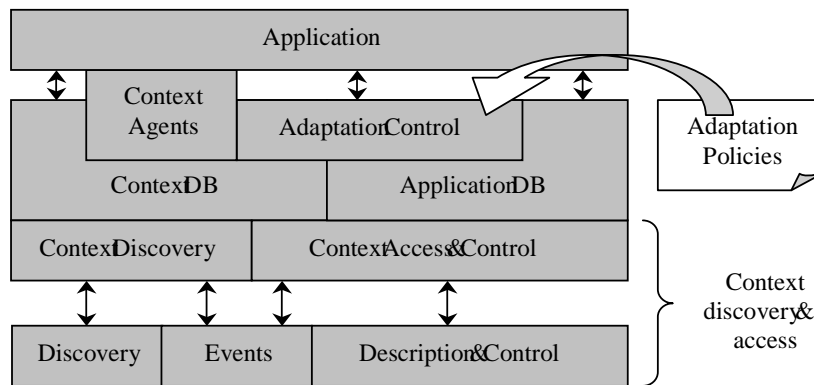


Fig2. The overall architecture of the system

All communications between the main components of the architecture are performed using HTTP as the transport protocol and XML to represent the format of the messages. The reason for this decision was based on the fact that these technologies support the lightweight integration of distributed components. Other alternatives, such as CORBA and RMI, are too heavyweight and would require additional infrastructure, e.g. ORB or RMI daemon respectively.

5.2 Context Discovery and Access

This part of our architecture is responsible for locating contextual information. These services can be either built into the system (e.g. a body temperature sensor) or services provided by other devices (e.g. active devices in the user's environment). In order to discover context services when they are available, a context-aware application must be able to discover context services when they are available. For this reason, a considerable part of the architecture is based on the UPnP architecture [19]. In more detail, the mechanism for discovering context services is based on services advertising themselves to interested context-aware applications.

Once a service is discovered, the platform retrieves the XML description of the device, an abbreviated example of which is shown in Figure 3. The XML description provides the access points for sending control messages and for event notifications. In addition, it defines the service and the types of information that the service can provide. Given this information, the platform is capable of fusing the information from the state of the service changes.

Context-aware applications provide services that provide context-aware information to the surrounding environment. In a mobile environment, it is important to be able to receive notifications when the architecture is based on the UPnP mechanism for discovering available services (using multicasting) to themselves.


```

<service>
  <category> Location </category>
  <type> GPS </type>
  <action>
    <name> getXCoord </name>
  </action>
</service>

```

Fig3. Sample of the XML definition of a service

5.3 Context Database

This component serves as a registry for all those context services currently available. An important part of the context database is a classification of the context services into context types. More specifically, each context service has to specify the type of context that it supports, e.g. both the GPS and the network-based locator provide location information. This specification of context type is achieved using an XML template that defines the kind of information this application can retrieve: the specific contextual data the service used for acquiring the data.

This method of hiding the actual mechanism for retrieving contextual information allows the platform to coordinate the access to context services. Moreover, it offers a platform with the potential to switch between different adaptation policies that have been specified.

More specifically, each context service has to specify the type of network-based locator provided. This specification of context type is achieved using an XML template that defines the kind of information this application can retrieve: the specific contextual data the service used for acquiring the data.

Moreover, it offers a platform with the potential to switch between different adaptation policies that have been specified.

5.4 Application Database

The application database serves as a repository for all applications running on the system. The application actually implementing the adaptive behaviour of this application developer is also responsible for exposing its behaviour to the platform by ensuring that the application registers all of its adaptation mechanisms with the application database.

The description of the adaptation mechanisms should be able to trigger this mechanism. This information is used to coordinate the triggering of the adaptation mechanisms. In figure 4, we illustrate an example of the information in XML concerning its adaptive function.

The adaptation mechanisms of all applications running on the system. The application actually implementing the adaptive behaviour of this application developer is also responsible for exposing its behaviour to the platform by ensuring that the application registers all of its adaptation mechanisms with the application database.

The description of the adaptation mechanisms should be able to trigger this mechanism. This information is used to coordinate the triggering of the adaptation mechanisms. In figure 4, we illustrate an example of the information in XML concerning its adaptive function.

5.5 Context Agents

A context agent is a piece of code that can be plugged into the platform in order to perform the application-specific manipulation of contextual information.

A context agent is a piece of code that can be plugged into the platform in order to perform the application-specific manipulation of contextual information.

```

<application>
  <name> WebBrowser </name>
  <adaptationMode>
    <name> lowBand </name>
    <trigger>
      <context> availableBand </context>
      <condition> lessThan-9600 </condition>
    </trigger>
  </adaptationMode>
</application>

```

Fig.4. Sample of the XML definition of application's operation modes. The XML based description provides the different operational modes of the application, coupled with the contextual trigger that would make the application switch into that mode.

A common case of context manipulation is the combination of primitive context information for constructing a complex type of contextual data. For example, a context agent plugged into the platform can combine location data and current time in order to provide location- and time tracking information similar to the data used by the Stick-Notes system [17].

To present even more clearly the operation of a context agent we will present an example based on the GUIDE system [2,3]. The GUIDE system is a mobile electronic context-aware tourist guide. As part of its functionality it provides information about tourist attractions in HTML format triggered by the location of the user. In order to introduce the GUIDE system into the platform we need to split the application into two parts: an ordinary web-browser and location-triggered HTTP proxy. The HTTP proxy operates as a context agent which is plugged into the platform and has direct access to the location information provided by the platform. Triggered by changes in location, the agent can request the appropriate HTML data from the content server.

The key motivation behind introducing the notion of context agents is to enable the development and distinguish the functionality of the application from the acquisition and manipulation of context data. Importantly, this allows context agents to be used for integrating non-context-aware applications (like an ordinary web browser) into a context-aware system.

5.6 Adaptation Control

This module is responsible for monitoring the status of contextual triggers and making decisions about the behaviour of the platform and the applications. The decision taking procedure is based on a set of adaptation policies specified by the user. These adaptation policies are specified by the application running on the system and among them the finishing priorities both among the sources of the system. This prioritisation represents the importance of these entities according to these needs.

The adaptation control is further divided into two sub-modules: *internal adaptation* and *external adaptation* explained below.

Internal adaptation. This module coordinates the context monitors that are required for all applications running on the system and on the platform itself. The adaptation actions that are acquired are tightly coupled with the context classification (in section 5.3). Context services are classified into context types according to the type of information they provide. For example, location mechanisms would be members of the same type of context, i.e. location. Both these mechanisms can provide similar types of information but have different specifications and different requirements. When the system gets into a state whereby one of the mechanisms is favoured (according to the adaptation policies specified by the user) then the platform will switch to the mechanism that is preferred (as illustrated in figure 5).




Fig5. Internal adaptation switching between different location mechanisms.

In the example discussed all applications that require location information do so via the location proxy object. Thus, the actual mechanism for retrieving contextual information is hidden from the application. This allows the platform to switch between mechanisms without affecting the applications involved. To illustrate this, consider the example described in section 2.1.4 and a scenario in which it is crucial to reduce aggregate device power consumption. In such a scenario, the platform could reduce power consumption by switching off the GPS device and using the network location in order to retrieve location information.

External adaptation. This module is responsible for coordinating the adaptive behaviour of the applications running on the system. Its operation is driven by both a set of adaptation policies defined by the user and the XML description of each application specifying its various operation modes (as described in section 5.4). Recall that the XML description of each operation mode is marked according to its effect on resource utilisation (e.g. power, network) and the use of context services. The adaptation control module can use this information in order to decide which operation mode should be triggered under each potential set of circumstances. In more detail, when resources (power, network, etc.) become unavailable the adaptation control module can use this information to decide which adaptation mechanism with the lowest resource requirements should be used. In order to clarify this approach, consider the following example (which is illustrated in figure 6). Two adaptive applications run on a mobile device: a web browser and a video player. The adaptation policies specified by the user define the priorities between the applications and the resource requirements (such as the power and network usage) of each (the greater the priority, the lower the resource requirements). The adaptation control module is aware of the adaptive modes that these applications can support by accessing the application database. It also knows the status of all the contextual data that is available on the system. When

any of the contextual triggers reach a value that triggers reaction by the system the adaptation control has to decide which adaptation mechanisms should be invoked. In the example presented, both power and available bandwidth are low. However, the adaptation control would choose to overcome the power problem, because the user has specified power as the most important resource. Not that enabling has set a specific priority as in this case enabled the platform to overcome an area of potential conflict (as highlighted in section 2.1.2).

gers reaction by the system the adaptation control has to decide which adaptation mechanisms should be invoked. In the example presented, both power and available bandwidth are low. However, the adaptation control would choose to overcome the power problem, because the user has specified power as the most important resource. Not that enabling has set a specific priority as in this case enabled the platform to overcome an area of potential conflict (as highlighted in section 2.1.2).

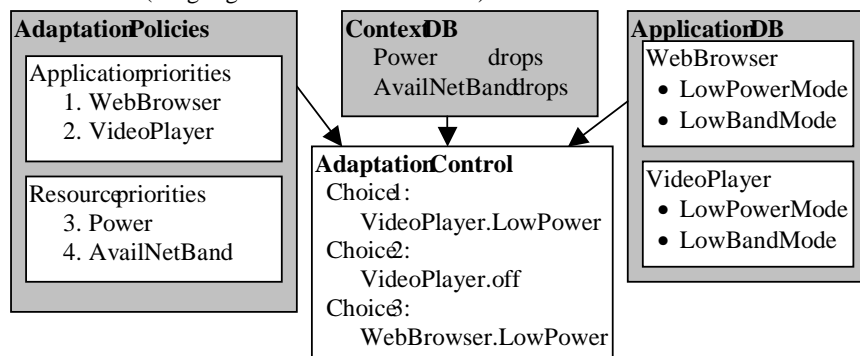


Fig.6. External adaptation making adaptation choices according adaptation policies

In order to decide which application to invoke the adaptation control component checks the prioritisation of the applications and picks the application that is less important. In this case (i.e. the video player). The adaptation control component then triggers the video player to switch to low power mode. If the power reduction resulting from this action is insufficient then the adaptation control component would proceed with the next course of action, i.e. turning off the video player and triggering the web browser to enter low

adaptation control component checks the application that is less important. In this case (i.e. the video player). The adaptation control component then triggers the video player to switch to low power mode. If the power reduction resulting from this action is insufficient then the adaptation control component would proceed with the next course of action, i.e. turning off the video player and triggering the web browser to enter low

6 Conclusions

In this paper, we have argued that existing architectural approaches for supporting adaptive mobile applications have a number of shortcomings. Furthermore, analysis of these shortcomings has led to the identification of a set of architectural requirements. Namely: support for a common contextual space, mechanisms to support co-ordinated adaptation between multiple adaptive applications and support for user defined adaptation policies. We have also described the architecture of our platform, which has been designed to meet these requirements and which enables mobile systems to extend their awareness of all relevant contexts that might affect overall system adaptation policies. Fundamental to our approach is the idea of having system-wide decision making policies that consider the most efficient adaptation outcomes from a number of possible adaptations. This is achieved by requiring the applications to provide information about themselves, their adaptation mechanisms, and the contextual triggers that affect their behaviour.

In this paper, we have argued that existing architectural approaches for supporting adaptive mobile applications have a number of shortcomings. Furthermore, analysis of these shortcomings has led to the identification of a set of architectural requirements. Namely: support for a common contextual space, mechanisms to support co-ordinated adaptation between multiple adaptive applications and support for user defined adaptation policies. We have also described the architecture of our platform, which has been designed to meet these requirements and which enables mobile systems to extend their awareness of all relevant contexts that might affect overall system adaptation policies. Fundamental to our approach is the idea of having system-wide decision making policies that consider the most efficient adaptation outcomes from a number of possible adaptations. This is achieved by requiring the applications to provide information about themselves, their adaptation mechanisms, and the contextual triggers that affect their behaviour.

References

1. Advanced Configuration and Power Interface Specification, Revision 1.0, Intel/Microsoft/Toshiba (1999).
2. Cheverst, K., N. Davies, K. Mitchell, A. Friday. Experiences of Developing and Deploying a Context-Aware Tourist Guide: The GUIDE Project. In: Proc. of MOBICOM'2000, Boston ACM Press (2000)
3. Davies N., K. Cheverst, K. Mitchell, A. Friday.: Caches in the Air: Disseminating Information in the Guide System. In Proc of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'99) (1999)
4. Davies N., A. Friday, S. Wade, G. Blair.: L²imbo: A Distributed Systems Platform for Mobile Computing. In ACM Mobile Networks and Applications (MONET) Special Issue on Protocol and Software Paradigms of Mobile Networks 3(2) (1998) 143-156
5. Dey A., Abowd G., Salber D.: A Context-Based Infrastructure for Smart Environments. In: Proc of the 2000 Conference on Human Factors in Computing Systems (2000)
6. Elis C.: The Case for Higher-Level Power Management. In Proc of HotOS (1999)
7. Flinn J., M. Satyanarayanan.: PowerScope: A Tool for Profiling the Energy Usage of Mobile Applications. In: Proc. of the Second IEEE Workshop on Mobile Computing Systems and Applications (1999)
8. Friday A., N. Davies, G. Blair, K. Cheverst. Developing Adaptive Applications: The MOST Experience. In Journal of Integrated Computer-Aided Engineering 6(2) 43-57
9. Golland, Y., Cai T., Leach P., Gu Y., Albright S.: Simple Service Discovery Protocol, Version 1.03. IETF Internet-Draft. <http://www.ietf.org/internet-drafts/draft-cai-ssdp-v1-03.txt>
10. Joseph A., J. Tauber, F. Kaashoek. Mobile Computing with the Rover Toolkit. In: IEEE Transactions on Computers Special Issue on Mobile Computing 43(3) (1997)
11. Katz R.: Adaptation and Mobility in Wireless Information Systems. In: IEEE Personal Communications 1(1) (1994) 6-17
12. Kravets R., P. Krishnan.: Application-Driven Power Management for Mobile Communication. In: Fourth ACM International Conference on Mobile Computing and Networking (MOBICOM'98) (1998)
13. Kunz T., J. Black. An Architecture for Adaptive Mobile Applications. In Proc of the 1st International Conference on Wireless Communications (Wireless'99) (1999)
14. Long, S., R. Kooper, G.D. Abowd, C.G. Atkeson.: Rapid Prototyping of Mobile Context-Aware Applications: The Cyberguide Case Study. In: Proc. of the 2nd ACM International Conference on Mobile Computing (MOBICOM) (1996)
15. Marzullo K., R. Cooper, M. Wood, K. Birman.: Tools for Distributed Application Management. In IEEE Computer 24(8) (1991) 42-51
16. Noble B., M. Satyanarayanan, D. Narayanan, J.E. Tilton, J. Flinn, K. Walker.: Agile Application-Aware Adaptation for Mobility. In: Proc of the 16th ACM Symposium on Operating System Principles (1997)
17. Pascoe J.: The Stick-Note Architecture Extending the Interface Beyond the User. In Proc. of the International Conference on Intelligent User Interfaces (1997)
18. Ferry D.B., M. Theimer, K. Petersen, A. J. Demers. Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System. In Proc of the 5th ACM Symposium on Operating System Principles (1995)
19. Universal Plug and Play Device Architecture, Version 0.91, Microsoft Corporation, March 2000 http://www.upnp.org/download/UPnP_Device_Architecture.mht
20. Weiser M.: Some Computer Science Issues in Ubiquitous Computing. In Communications of the ACM 6(7) (1993) 75-84