Active Networking based Service Infrastructure for Cellular WLANs

Reiner N. Schmid[1], Cornel Klein[1], and Thomas Becker[2]

[1] Siemens AG, Munich, Germany

[2] Fraunhofer Institute FOKUS, Berlin, Germany

Abstract. An approach of managing a service infrastructure for wireless local area networks (WLANs) is presented. This approach uses Active Network technology as developed in the IST project FAIN in order to deploy and run services at appropriate locations in the network. By introducing a service which monitors traffic and adaptively controls the access of mobile users it is possible to dynamically bind particular users to specific access points, thus transforming conventional WLAN to cellular WLAN. This service is made up from components in the data and control plane of the network. Additional components in the management plane are used for the deployment and runtime management of the service components. The usefulness of this approach is corroborated with the outline of key scenarios.

1 Introduction

WLAN [2] networks have become very popular for a wide range of users and services and are rapidly deployed as public hot spots in areas such as airports, conference rooms, restaurants, shopping malls and public fairs. WLAN has emerged as the technology of choice used for wireless data transfer in public areas. From a technological point of view it is comparable to Ethernet as it establishes a shared medium for data exchange without provisions for exclusive access. This concept is quite opposite to mobile communication technologies such as GSM or UMTS, which have mechanisms for exclusive channels and cellular structures, being the basis for the provisioning of Quality of Service (QoS) guarantees and other mechanisms for resource control.

WLAN, as standardised by IEEE, offers only a limited support for advanced cellular mechanisms and roaming. Though there is emerging support for inter access point communication in the recently established standard IEEE 802.11.f – defining communication between the access points about re-association of terminals – this is limited to facilitate the deployment of large scale networks consisting of many WLAN access points.

WLAN, based on IEEE 802.11b, allows for transmission rates up to 11 Mbps, which is sufficient for audio transmission and to a limited extent also for video streaming given that the streaming bandwidth is adapted to the available transmission bandwidth (see e.g. [3]). Though the newer and faster WLAN standards (e.g. IEEE 802.11g) allowing

N. Schmid R., Klein C. and Becker T. (2004).
Active Networking based Service Infrastructure for Cellular WLANs.
In *Proceedings of the 3rd International Workshop on Wireless Information Systems*, pages 134-141
DOI: 10.5220/0002674801340141

for higher data rates are now available, support for dedicated Quality of Service is limited and not sufficient to support service environments with different services using the same WLAN network at the same time.

In this paper we focus on WLAN application scenarios in semi-public areas (airports, hotels etc.), which require a flexible service provisioning infrastructure concept in order to fulfil the requirements of users and services alike. With respect to these scenarios as discussed in section 2, more flexible and sophisticated mechanisms for support of service infrastructure by WLAN are required. Though this isn't part of the MAC layer specification of WLAN by IEEE, it is shown in this paper that this service infrastructure can be implemented by using Active Networking [1], a promising method of solving the demands of next generation networking infrastructures. The basic concept of Active Networks is to allow third parties (service providers, end-users, etc.) to inject application specific service code into the network. Active Network nodes expose open interfaces and execution environments for running services, thereby making the network 'programmable'. Wireless local area networks (WLANs) prove to be a particular interesting application domain for Active Networking [4].

Active Networking concepts, in particular the Active Node architecture of FAIN as presented in section 3, are used to implement a flexible, application adaptable resource management of WLAN networks. This includes the establishment of virtual service environments and dynamic access control for WLAN resources thus providing cellular mechanisms for WLAN management.

The main idea of our approach is to transfer concepts of cellular mobile networks to WLAN establishing a cellular WLAN, thus complementing out-of-the-box WLAN with state-of-the-art concepts of mobile networks. This is shown in Section 4 which presents the details of a service for dynamic control of access to a WLAN. Conclusions are discussed in Section 5.

2 Mobile Active Networking Scenarios

We mainly consider WLANs in so-called *semi-public areas*. By a semi-public area we denote a geographically separated area with a network infrastructure which is populated by a huge variety of stakeholders, but which is owned and controlled by a particular institution, e.g. a company. In contrast to completely private areas with very restricted access, it is frequented by a large number of users with different objectives and interests. Some examples for such semi-public areas and the users of their WLAN infrastructure are:

- airports with users like passengers, airlines, shops, airport personnel;
- trade shows with users like visitors, exhibitors, organiser's personnel;
- shopping centres providing WLAN support for visitors, shops clerks, service personnel, etc.;
- hotels with guests, conference organiser, hotel personnel;
- corporate campuses with employees, external visitors, different organisational units.

In all these cases there are a couple of groups with different requirements when using the WLAN infrastructure. For instance, at a trade show the visitors may prefer best-effort or QoS-based Internet access, depending on their willingness to pay. Exhibitors

may want to use WLAN for providing location-based services to visitors, e.g. to attract people to visit their exhibition places. Additionally, exhibitors could use WLAN for their internal communication as do the employees of the trade show organiser. In that case it is likely that WLAN based telephony will only be used if it is ensured that it doesn't interfere with other WLAN usage, e.g. ongoing data transfer.

To summarise, the requirements of the different users towards WLAN operation are quite specific and diverse. Ad-hoc approaches, where different WLANs are installed and operated by (or for) the different stakeholders aren't suited to accommodate their joint requirements for the following reasons:

- Setting up separate WLANs is likely to be more expensive than using a shared infrastructure, in particular, if they are used only for a short period of time.
- Overall Quality of Services requirements can only be implemented if access to the underlying shared resources as WLAN channels and bandwidth is controlled by a common entity.
- Moreover, a powerful shared infrastructure is beneficial for high-volume peak times, as, for instance, an application with high bandwidth requirements can be shifted to a specific WLAN channel or access point.

Since the involved stakeholders and their requirements in the presented examples are changing quite frequently, static approaches to the configuration of such a common, shared infrastructure aren't well suited. In particular, it is rather unlikely that all required services are configured by the infrastructure's operator in advance. What we need is an infrastructure which among others implements the following requirements:

- *Delegation*: part of the management of the infrastructure can be delegated to its stakeholders, allowing them to manage their part of the network on their own. The provided management functionality should be as flexible as possible w.r.t. to the kind of applications they have in mind.
- *Virtualisation*: the different stakeholders shouldn't interfere with each other, but should get their own (virtual) private network, both w.r.t. to management functionality and user traffic.
- *Resource Management*: it should be possible to flexibly allocate, release and monitor the underlying shared resources on a per stakeholder/per users basis.

The Active Networks based service infrastructure presented in this paper realises these requirements complementing existing approaches for deploying and operating large WLANs. For this reason we don't deal with issues like network planning (e.g. access point placement), but we rather take for granted viable approaches for dealing with these problems.

3 FAIN Active Node

The management layer of the FAIN Active Node [4] is of particular importance in the FAIN architecture. The purpose of this layer is to manage active services in the scope of a network node on the data, control, and management planes. It serves as an abstraction layer between services and the node's operating system by abstracting from operating

system specific interfaces and making them available as so called *basic node services* to higher-level services. In this way the management layer provides a uniform way to deal with functionality provided by the node as well as dynamically added extensions in the form of active services.

Since services may be implemented using different technology and thus require different execution environments (EEs) the FAIN node architecture doesn't define a single type of EE but rather allows for a variety of them. For example, complex services in the control or management plane may be well executed in a user-space EE providing considerable flexibility and robustness while services in the data plane may require a kernel-space EE for performance reasons.

We want to mention two types of EEs developed in the FAIN project, namely the JAVA-based EE [4] and the PromethOS EE [5]. With the JAVA EE complex services can be rapidly prototyped making use of the broad functionality provided by the JAVA API. It is also used for the implementation of the FAIN active node's management layer. The PromethOS EE is an extension of the Linux Netfilter framework and running in kernel-space. In particular it allows to dynamically add services in the data path in the form of modules. Thus it is predestined for services working on data packets and requiring high-performance which wouldn't be possible when data had to be copied from kernel-space to user-space and vice versa.

For specifying, implementing, and managing services a component-based approach was chosen. This approach offers a couple of advantages:

- Aspects such as lifecycle management, configuration, access control, monitoring, etc. can be separated from the service logic and are supported by the components' runtime environment. Thus the developer of a service can concentrate on the service logic and frequently needed aspects don't have to be implemented over and over again.
- Complex services can be built by combining available components.
- By providing means to dynamically interconnect components a high degree of flexibility can be reached. After the initial setup services may be reconfigured during runtime by adding new components, removing components, or changing connections between components.

Other systems with different or no component abstractions can be wrapped with proxy components residing in the management layer thus making available the mentioned advantages. Similarly, the abstraction of modules provided by the high-performance PromethOS EE was mapped to components on the management layer.

The FAIN Active Node architecture defines some basic node services which are available to later installed services. The basic services comprise management of component environments, control of access to services and resources, packet dispatching among components, and traffic control.

There are two different kinds of component environments: the already mentioned *execution environments* (EEs) implemented with a specific technology and *virtual environments* (VEs) which abstract from a particular EE implementation and are used to group resources belonging to the owners of VEs. A special *privileged VE* is owned by the node's operator and has full access to the node's resources. This is where the basic node services are installed when the node is booted.

The next section will describe a specific service used for controlling the access to a WLAN. The service is made up from components in the data and control planes and is running in two different types of execution environments.

4 Wireless Access

Based on the component model of the FAIN project a service for controlling access to a WLAN, that takes into account the requirements posed by scenarios as discussed in section 2, has been implemented and demonstrated during the FAIN project [6]. The service effects load balancing between a number of wireless access points. In order to provide support not only for this particular service, but for a variety of services – a key feature in this application domain – a generic framework for access control has been developed. This framework makes use of the advantages of the kernel-space PromethOS EE for processing data packets and the user-space JAVA EE for control purposes. With this framework one can build access controllers equivalent to commercially available solutions. Moreover, by exploiting the flexibility provided by the EEs it is possible to deploy new services on-the-fly and to customise their behaviour according to different applications' needs. The establishment of application dependent resource control and virtual network environments for the provision of resource shares to different users are the key features of our implementation.

The framework for access control to wireless networks defines three entities situated at different locations:

- terminal daemons located at the end users' terminals,
- a WLAN access controller (WAC) located on a potential active node in the access network, and
- components in the data plane on one or more active nodes supporting the WAC.

A sample network setup showing the deployment of the different components is presented in figure 1. The details of their functionality are explained in the subsequent paragraphs.

Terminal Daemon: The terminal daemon is responsible for managing the WLAN driver of the terminal. It determines the access points to be used for association and controls the WLAN roaming behaviour. Further, it obtains information about the available access points and their signal quality from the WLAN driver. The terminal daemon sends this information to the WAC and receives instructions which are then enforced using the terminal's WLAN driver.

WLAN Access Controller: The WAC's main task is to exercise control on all issues related with the management of users and usage of the WLAN. Based on the information received by the terminals and its own monitoring results, it decides on access and load distribution in the network. Configuring the FAIN active nodes is a further task of the WAC. The WLAN Access Controller is implemented based on the JAVA EE of FAIN.

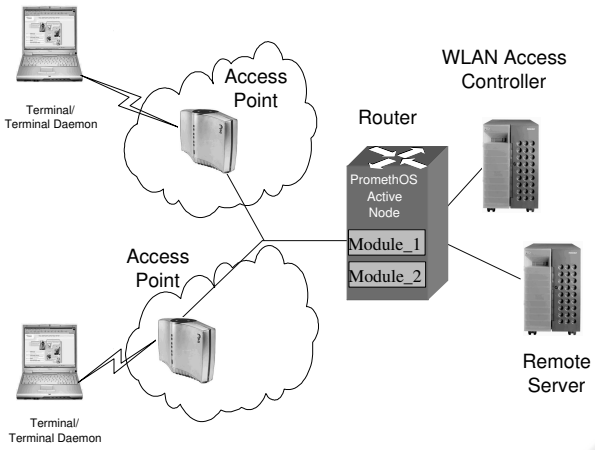

Fig. 1. Deployment of WAC system components

PromethOS Modules: The components in the data plane EE – i.e. modules in the PromethOS EE – are responsible for IP layer decisions concerning the traffic in the WLAN network. They are configured by the WAC. In this respect the functionality of the network node may be freely configured by the deployment of PromethOS modules. However, a certain minimal set of basic functionality is always required:

- routing/bridging between different subnets of WLAN access points and other networks,
- load monitoring of data flows from and to access points per user or by other monitoring criteria,
- routing to simulate bridging functionality of the active node, and
- receiving messages destined for the WAC.

An example demonstrating the interaction of the different entities for a load balancing service is shown in Figure 2. It shows a handover between different WLAN access points. The terminal daemons are constantly doing measurements about the available access points and their signal strengths. This measurements are reported to the WAC. Similarly, the traffic through the active node is measured and results are reported to the WAC. Based on this information the WAC decides which terminal should be assigned to which access point. In the case that a terminal should switch to another access point the WAC sends instructions to the corresponding terminal daemon.

Based on the FAIN component model a very rapid development of the different components could be made. In particular the JAVA components can be very efficiently programmed. The PromethOS EE components are more sophisticated to program, as they are implemented as Linux kernel modules.

Depending on the various needs of applications running on mobile terminals different services can be deployed on demand. Making use of the component-based approach a new service for wireless access control is deployed by installation of modules for the

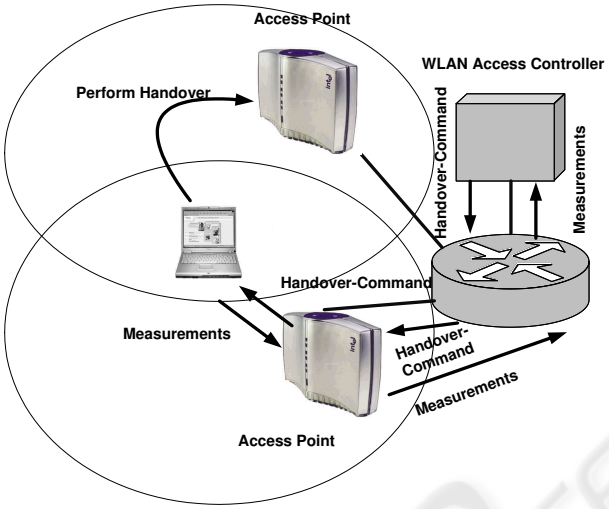

Fig. 2. Interaction of entities in a handover scenario

PromethOS EE and WAC components for the user-space EE. The installation of service components to different EEs is coordinated by the VE to which the EEs are attached. The VE concept with the separation of resources allows to install multiple services in parallel sharing the same underlying infrastructure.

5 Conclusions

In this paper, we presented the FAIN architecture for managing active nodes and an application based on it for wireless networks. The introduction of virtual environments in FAIN allowed integrating several execution environments with potential different implementation technologies. Further, physical resources could be partitioned among several node users – the service providers – with the help of virtual environments. To achieve a flexible and fine grained control over service deployment and management a component-based approach was chosen for the node level management layer. The application in a wireless access systems turned out to offer a considerably increased flexibility compared to commercial off-the-shelf solutions already available for access control in WLAN networks. Another strength of the Active Networking concept is the possibility to effect an application dependent behaviour of network elements. The potential benefits of these capabilities that are easily installed in the framework for wireless access presented here remain to be investigated and are a field of further study.

Acknowledgement

This paper describes work undertaken in the context of the FAIN project (www.ist-fain.org). We would like to thank our colleagues of the FAIN consortium for valuable

discussions, in particular Lukas Ruf from ETH Zurich for implementing and supporting the PromethOS EE. The FAIN project was partially funded by the Commission of the European Union as IST project 10561.

References

1. Bhattacharjee, S.: Active networks: Architectures, Composition, and Applications. Ph.D. Thesis. Georgia Tech (1999)
2. IEEE Standard 802.11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications (1999).
3. Kassler, A., Schorr, A.: Generic QoS aware Media Stream Transcoding and Adaptation. In: 13th Packet Video Workshop. Nantes, France (2003)
4. A. Galis, S. Denazis, C. Brou, and C. Klein: Programmable IP Networks: Management and Rapid Service Deployment, Artech House London (2004).
5. Keller, R., Ruf, L., Guindehi, A., Plattner, B.: PromethOS: A Dynamically Extensible Router Architecture Supporting Explicit Routing. In: Proceedings of the Fourth Annual International Working Conference on Active Networks (IWAN 2002). Lecture Notes in Computer Science, Vol. 2546. Springer Verlag, Berlin Heidelberg New York (2002)
6. Flury, P. (ed.): FAIN Demonstrators. FAIN Public Deliverable 40 (2003). http://www.ist-fain.org/deliverables/del40/d40.pdf