The Composition Agreement in Ambient Networks

Rody Kersten

December 2006

Contents

Chapter 1

Introduction

Over the past five decades we have witnessed the introduction of the computer into our everyday lives. This is a trend that keeps on evolving in an extremely quick way. The main reason for this is the fact that computing power doubles every 1,5 years, as Gordon E. Moore already predicted in 1965. His exact law, which holds to this day, states that "the number of transistors on integrated circuits, which is a rough measure of computer processing power, doubles every 18 months" [1]. Not only does this mean that the power of computers increases exponentially, it also means that computers can become smaller and smaller. This, together with a similar decrease in expense, is why we are surrounded by more and more computers in our everyday lives. Nowadays almost every electronic device is equipped with some kind of a computer, and we are already carrying at least one, quite powerful, computer on our body: our cell-phone. This gives us all the reason to believe that in the near future, almost every new product will be equipped with a computer if a use for it.

One obvious use for these computers, useful to almost every artifact, is communication. For example, if your refrigerator has a computer, which is connected to the Internet, it will be able to order a new piece of cheese for you, as soon as you finish the old one. The ideal here would be to enable all these devices to communicate with each other, share services and resources, and to simply be aware of their environment.

A co-effect of these developments is the fact that we are getting less aware of the computers around us, simply because they are everywhere and perform all sorts of tasks for us that we cannot even remember performing ourselves. The time in which sitting behind a desktop PC was the only computer interaction we got is now far behind us. What we are heading towards is what can be referred to as *Ambient Intelligence*, *Ubiquitous Computing* or *Pervasive Computing*: A seamless environment of computing, advanced networking technology and specific interfaces [2].

1.1 Ambient Intelligence

Ambient Intelligence (AmI) is a vision in which humans will be surrounded by intelligent interfaces supported by computing and networking technology that is embedded in everyday objects such as furniture, clothes, vehicles, roads and smart materials - even particles of decorative substances like paint [2]. Most of this technology will be seamless to humans, hence the name *Ambient* Intelligence; it will be in the ambiance of the human.

An AmI environment should be:

- Unobtrusive. Humans should not be too aware of the technology involved in AmI. Interaction with the AmI environment should be relaxing and enjoyable.
- Aware. The AmI environment should be aware of human presence and adopt to certain needs and desires.
- Trust-worthy. Lots of precautions will have to be taken in order for AmI not to invade the privacy of its users.
- Controllable. The On/Off switch should be within reach. A great danger of AmI is that people may get the feeling that AmI is controlling them, instead of them controlling AmI.

The basic idea of AmI is that computers are everywhere. People will walk around with some kind of wearable computer, worn like a watch around the wrist, or as a chip, implemented into the arm. This personal computer will be able to make all sorts of connections to all sorts of heterogeneous networks, for example a car's VAN (Vehicle Area Network). The car's VAN, for its part, will be able to connect to other cars' VANs and be sure not to crash into them. It also will be able to connect to a central database, tell it where it is heading and how many more people can fit in the car. Other people will be able to check this database, so maybe we can pick up some carpoolers, who's wrist-computer will make an arrangement for them, and then, after driving with us, automatically pay for the costs.

These are just a few of the many possible uses for Ambient Intelligence. Other uses include all sorts of communication with other people, adjusting different environments to your needs and wishes (for example a hotel room), assistance in education and assistance in basically every daily task. Some nice scenarios for this use can be found in [3].

For AmI to happen, of course, a lot of technologies must be (further) developed. The key technological requirements are, according to [3]:

- Requirement 1: Very unobtrusive hardware
- Requirement 2: A seamless mobile/fixed communications infrastructure
- Requirement 3: Dynamic and massively distributed device networks
- Requirement 4: Natural feeling human interfaces
- Requirement 5: Dependability and security

Requirement 3 is where Ambient Networks fit into this vision. As seen from the above example use for AmI, there's a great need for advanced networking. Because of the mobility of computers, ad hoc networks need to be formed constantly, network composition has to be almost instant and services and resources need to be easily shared. As I will show in the next chapter, Ambient Networks achieve this task.

1.2 Fourth Generation Wireless Communication

The term Fourth Generation (4G) is used to describe a spectrum of wireless communication technologies. Its predecessor, 3G, is currently being implemented and involves techniques like UMTS. 2G involved GSM, among others.

One of the objectives of 4G is that it will allow entities to connect to the Internet and each other through a variety of devices and standards anytime, anywhere, and at a wide range of speeds, from narrow band to broadband [4].

The Ambient Networks project will fulfill this need for increased connectivity and portability. Its main aim is to develop "an architecture for communication networks beyond 3G" [5].

1.3 Paper Outline

This paper describes my work for my bachelor thesis. The research question was: *How can Ambient Networks agree on composition terms?*

In this paper I try to explain exactly what these Ambient Networks are and how they work. I will introduce Ambient Networks in the next chapter. This chapter is the result of the literature study I have conducted on Ambient Networks in general. In chapter 3 I will describe the results of my literature study on the Composition Agreement. This is the agreement that is negotiated by composing Ambient Networks, which determines the ontology of the composed network. I also try to gain more insight in the principle of network composition by formalizing the negotiations that take place between Ambient Networks before composition and building a prototype application of these negotiations. This also serves to reveal potential weaknesses in the design of these negotiations and to test the overall feasibility of the concept and will also be described in chapter 3. In the final chapter I will draw conclusions from my findings.

Chapter 2

Ambient Networks

The need for inter-connectivity between computers is increasing rapidly. More and more devices are being connected to each other and the Internet. These connections are also more often wireless and mobile. The nature of these connections, however, still differs a lot, which is very impractical to its users, be it end-users or application programmers. A whole set of different techniques apply to a wireless Bluetooth connection to a mobile phone, than to a wired UTP connection or even a direct USB cable link.

In fact, end-users don't need to be bothered with all these details. Even more so, Ambient Intelligence will only happen if network applications can be easily and quickly developed. For an application developer or an end user, it is of no interest at all how the underlying connection that he or she is using exactly works, as long as its there and is functioning.

The Ambient Network architecture is an abstraction from all these heterogeneous connections and protocols, which serves to simplify mobile, ad hoc networking, in a way in which it appears as homogeneous to the potential users of network services.

This view presents a method for abstraction from intensely varying network technologies and covering them by a uniform structure and interface.

In this chapter I will describe the findings of my literature study on Ambient Networks. I will first describe the Ambient Networks project, which is developing Ambient Networks in section 2.1. I will then introduce the AN design principles in section 2.2. Finally, I will explain the actual design in section 2.3.

2.1 The Ambient Networks Project

Ambient Networks is an integrated project by a collaboration of companies and universities like Ericsson, Nokia, Siemens, TNO and the technical university of Berlin, co-sponsored by the European Commission under the Information Society Technology (IST) priority under the 6th Framework Programme. The project addresses the strategic objective of "Mobile and Wireless Systems Beyond 3G".

The Ambient Networks project is currently in phase two out of three. Phase one was *Concept Exploration*, phase two is *System Engineering* and phase three will be *Wireless World Initiative Synthesis*. No articles have been published by the Ambient Networks project group since the end of phase one, so my research will concentrate on this first phase.

This means that all the techniques and systems presented herein, are all just concepts for now. This is the reason I will also build a prototype of the Composition Agreement negotiations; the concepts have not been proved feasible yet (at least to the public).

2.2 Design Principles for Ambient Networks

According to [5], the AN concepts are based on the following principles:

- **Principle 1: Ambient Networks build upon open connectivity and open networking functions.**

  One basically new way of defining networking in Ambient Networks is to remove architectural restriction on who or what can connect to what. In this open network structure, there are no end nodes, just networks. A computer or network does not need to know, and cannot know, whether it is connected to a computer or a network.

  This view provides an abstraction from the actual structure of the networks. The goal is to enable all networking services for connected networks, instead of just two connected nodes. Since there are no end nodes in this view, we talk about *end environments* rather than end nodes.

- **Principle 2: Ambient Networks are based on self-composition and self-management.**

  Because of the large number of ad hoc connections between networks, a need for human interaction in setting up these connections is not desirable. Ambient Networks must be able to manage themselves, their services and their connections. No preconfiguration or offline negotiation must be needed.

- **Principle 3: Ambient Network functions can be added to existing networks.**

  Ambient Networks have to be able to cooperate with existing networking technologies. If Ambient Networks are not compatible with existing technologies, they will never be implemented.

2.3 Ambient Networks Design

One of the most innovative ideas of Ambient Networking is the fact that there are no end nodes, just networks. So, the building blocks for larger network structures are networks, not PCs. Two or more of these networks (which may in fact be just a single node) can merge, and form one, bigger, Ambient Network. This process is called *Composition* and is described in more detail in section 2.3.1 and chapter 3. The current Internet, in contrast, has a pure end-to-end structure.

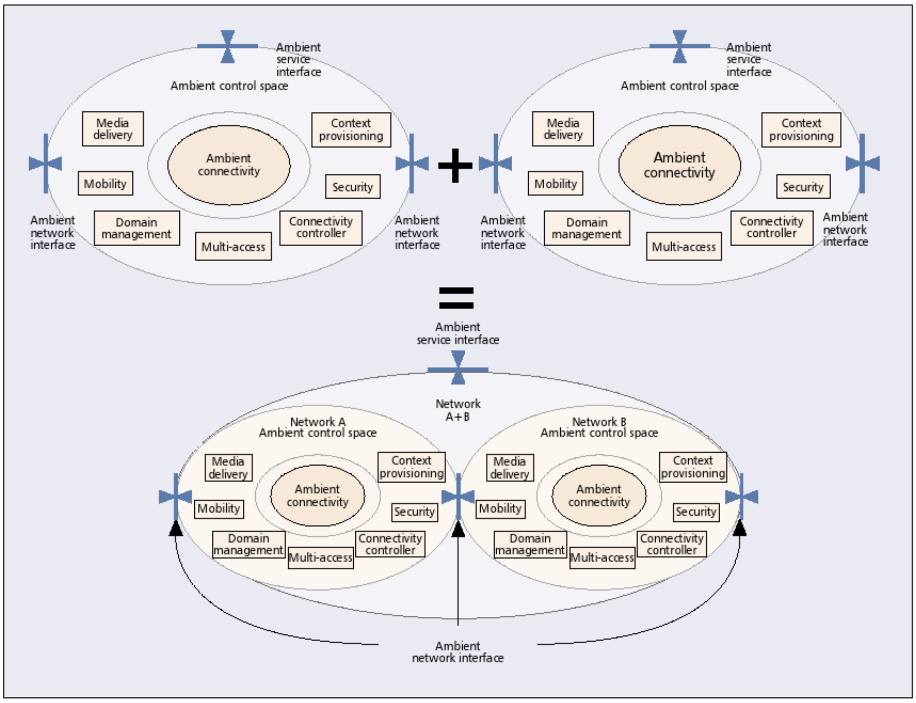

Figure 2.1: Partial Network Composition [5]

This idea of composition is the driving force behind AN. It is the reason that ANs can merge quickly and autonomously. An example of this composition is shown in figure 2.1. In this schema, the two networks merge and form a new *composed* network.

The key element in AN is the *Ambient Control Space* (ACS). The ACS contains all control functions. Examples of these control functions are mobility support, support for multi-access networks and a security function. "The ACS together with a (possibly legacy) connectivity network is called an Ambient Network" [5]. The Ambient Control Space is thoroughly covered in section 2.3.2.

An AN has three interfaces, the *Ambient Network Interface* (ANI), the *Ambient Service Interface* (ASI) and the *Ambient Resource Interface* (ARI), which is the interface to the underlying network connectivity layer. These interfaces are described in section 2.3.3

2.3.1 Ambient Network Composition

In an Ambient Network Composition two or more Ambient Networks are merging and forming an ACS together. All communications go through the ANI.

There are three kinds of compositions [6]:

- **Full composition** or **Network Integration**
  In a full composition or network integration the composing networks merge fully. The new network contains and controls all resources of all its members. The constituent networks give up all control.
- **No-composition** or **Network Interworking**
  Network interworking or no-composition can also be described as an attachment. This is how ANs connect to legacy networks. Basically, there is no composition; the networks just connect to each other, like in present-day networks.
- **Partial composition** or **Control Sharing**
  All other compositions can be described by the term partial composition. In a partial composition the composing networks share a subset of their resources and form a new ACS together. This is the type of composition depicted in figure 2.1.

Our focus will be on partial composition, because this is both the typical type of composition and the most interesting and innovative. According to [12] this is expected to be the most common case of composition.

To compose, the merging networks have to form a *Composition Agreement* (CA). This agreement contains all the conditions for composition the networks have agreed on. For example, a Composition Agreement can contain a deal in which Ambient Network A shares its Internet connection with Ambient Network B and AN B shares some files, a printer and maybe a LCD display with AN A.

The forming of this Composition Agreement is the main subject of this thesis. It will be discussed in great detail in chapter 3.

The process of composition can be best described by some examples. I will give a scenario for all three types of composition, starting with a partial composition. In these examples I will use situation specific names for each phase and omit the *Media Sense* phase. The Composition Process will be described in its general form in section 3.1.

This first scenario describes the meeting of Pete and Mary and the partial composition of their Personal Area Networks (PANs). A PAN is a computer network that is centered around you, which contains for instance your laptop, PDA and GSM.

Pete and Mary are having a business meeting at the office of the company Mary works for. Their PANs have a direct radio connection. Pete has some documents on his PDA which are open to all his business relations (his business card, some information about his company, et cetera). Mary's laptop has a connection to the Internet through the office intranet, which she is willing to share with any other people inside the office.

- **Phase 1: AN Discovery**
  ANs will have a discovery mechanism which will detect other ANs available for composition. The networks contact each other, based on this discovery mechanism, which makes both networks detect another network available for partial composition. They also determine that no other AN is involved in the possible composition.
- **Phase 2: Authentication and Authorization**
  The Authentication and Authorization procedure is carried out. What this encapsulates precisely can vary in different situations and types of compositions, because of different needs for security and identification. In this case the networks perform authentication by sending each other a certificate issued by a third party, like in SSL.
- **Phase 3: Composition Agreement Establishment**
  The ANs negotiate about the specifications of the composition. A Composition Agreement (CA) needs to be formed, which will determine the ontology of the composed network and which services will be shared. Pete will share his documents and will be able to access the Internet through Mary's laptop and the office intranet.
- **Phase 4: CA Realization**
  If a CA could be established, the networks now compose according to this CA. This new composed AN can be seen as a *virtual* network maintained by Pete and Mary's PANs. If necessary the composed network can hide the AN identity, ANI and ACS of its constituents.

The next example describes a no-composition. In this example, Pete is returning to his own office. Pete needs access to the Internet, so his PAN is constantly looking for ANs offering this service. On his way, Pete moves through an area covered by an infrastructure network. In this case, this is a network through which one can *rent* services, like Internet access.

- **Phase 1: AN Discovery**
  Pete's PAN and the infrastructure AN find each other using the discovery mechanism and determine that no other AN is involved in the possible composition.
- **Phase 2: Authentication and Authorization**
  The Authentication and Authorization procedure is carried out.
- **Phase 3: Service Level Agreement Establishment** Since this is a *no-composition*, there is no CA. Instead, a *Service Level Agreement* (SLA) is formed. This is an agreement between the provider and recipient of a service. SLAs are not AN specific and are in face already in use today. In

this case, the SLA states that Pete will get Internet access, at the expense of 10 cents per minute.

- **Phase 4: Attachment**
  If a SLA could be agreed upon, Pete's PAN can now attach to the infrastructure AN. There is no actual composition between the networks, thus no new ACS is formed.

The final example describes a full composition. In this scenario, Pete arrives at his office where he turns on his PC. Pete's PC will now fully compose with his PAN.

- **Phase 1: AN Discovery**
  Pete's PC's ACS and his PAN's detect each other using the discovery mechanism and determine that no other AN is involved in the possible composition.
- **Phase 2: Authentication and Authorization**
  The Authentication and Authorization procedure is carried out.
- **Phase 3: Composition Agreement Establishment**
  A Composition Agreement is formed which states a full composition.
- **Phase 4: CA Realization**
  The ANs will now fully compose. Both networks hand over all control to the newly formed ACS.

Figure 2.2 depicts these three examples. In the full composition the constituent networks are absorbed entirely by the composed AN. In the no-composition no new ACS is formed, only a link is created. In the partial composition the two constituent networks share a new ACS together.

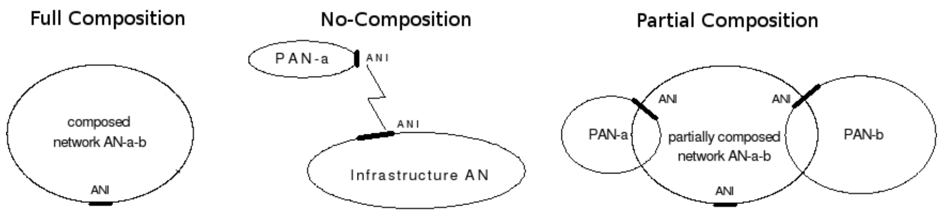

Figure 2.2: The three types of AN composition

2.3.2 The Ambient Control Space

The Ambient Control Space is the technical inside of an Ambient Network. It holds the set of control functions implemented in the AN. It is composed of multiple Functional Areas (FAs). Among others:

- **Composition Functional Area (C-FA)**
- **Quality of Service Functional Area (QoS-FA)**
- **Mobility Functional Area**
- **Congestion Control Functional Area (CC-FA)**
- **Security Functional Area**

These Functional Areas can communicate through the ANI, which makes cooperation between FAs of different ANs possible. The FAs of higher level Ambient Networks can be implemented in two ways: *centralized* or *decentralized*. In the centralized option the FA is actually implemented in the AN it belongs to. In the decentralized option the FA is implemented in one of the constituent networks of the AN and the communications to this FA are tunneled to this network.

Most of these FAs are optional. An AN implements a subset of the ACS functions.

Quality of Service will be a very important function in ANs, because of the large diversity in connections and, especially, the quality of wireless connections. Networks must be able to make a well-considered decision of which connections to use in certain situations.

Figure 2.3 depicts the Ambient Control Space and some of its Functional Areas and interfaces.

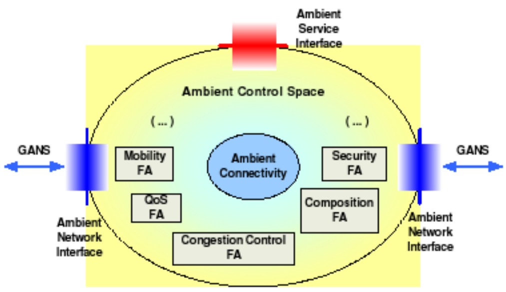

Figure 2.3: The Ambient Control Space

2.3.3 Ambient Networking Interfaces

An Ambient Network has three interfaces:

- **The Ambient Network Interface**
  The *Ambient Network Interface* (ANI) is used for all communications between the FAs of different ANs. "It offers standardized means to connect the functions of an ACS with functions of another domain" [5]. It can be

used for communication between either composed or non-composed networks. Two separate networks can use it to negotiate composition and a composed AN uses it on the inside for cooperation between FAs of all its constituents. It can be used for ACS signalling, information exchange and the discovery of other ANs.

A special protocol for communication over the ANI has been designed: *Generic Ambient Network Signalling* (GANS). This protocol is used for the exchange of information currently not sufficiently covered by existing protocols.

The ANI consists of multiple logical channels which handle either signalling associated with individual control space functions, or shared functionality.

- **The Ambient Service Interface**
The *Ambient Service Interface* (ASI) is used for accessing the services of an AN. It is the *upper-layer* interface to the control space [10].

- **The Ambient Resource Interface**
The *Ambient Resource Interface* (ARI) is the interface between the ACS and the underlying connectivity layers. The ARI specifies an abstraction from this connectivity and presents a uniform interface to the ACS FAs.

All three of these interfaces have a common structure that consists of some basic interface functions and some functions specific to the ACS implementation.

Dealing with an AN, external entities are not familiar with its ontology. "Even in a composed AN, only a single homogeneous control space is visible" [5]. An external entity connected to an AN will always find the same interfaces, regardless of the AN's ontology. This provides a great source for abstraction and scalability.

Chapter 3

The Composition Agreement

When two Ambient Networks compose, they do so according to the Composition Agreement (CA). This Composition Agreement contains the control sharing specifics to which the composing networks both agreed and determines the ontology of the composed network.

This composability will go far beyond what the Internet and mobile networks provide today. Network interworking shall not only be able at the level of addressing and package forwarding, it will also incorporate support for higher level sharing of services and resources. Composition will serve as a platform for the building of large networks, using uniform communication methods and deploying the resources of individual networks.

This chapter contains the results of my literature study on the Composition Agreement in the first two sections, my formalization of the negotiation process in section 3.3 and a description of the negotiation prototype I have built in section 3.4. In the final section I will discuss advantages and disadvantages of the centralized and distributed approaches to the CA negotiations.

3.1 The Composition Process

The process of AN composition has five phases: Media Sense, Discovery / Advertisement, Security and Internetworking Establishment, Composition Agreement Negotiation and Composition Agreement Realization. These phases are depicted in figure 3.1. This is not necessarily a one-way process, sometimes one or more steps back may also be taken. For example, after the negotiations, another service might be advertised, which remained hidden until some Quality of Service aspects had been agreed upon, which are a precondition for using this service. These steps are not shown in the figure because they can be taken from and to every phase in the model, so this would descrease its readablity.

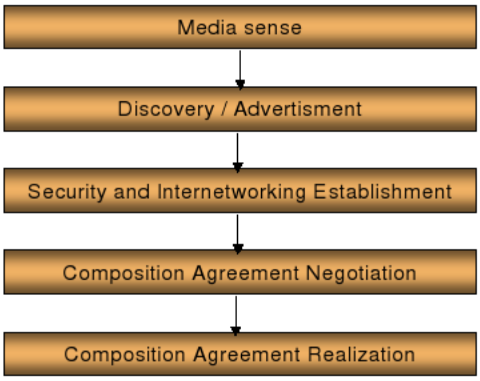

Figure 3.1: The five phases of the Composition Process [14]

The phases in the composition process will be described in detail in the next five sections.

3.1.1 Media Sense

In the Media Sense phase, an AN is searching for media for possible communication with other ANs. Media Sense may be triggered by different types of events, for instance [14]:

- A human network operator connects a new access point to its network.
- Two networks, managed by a human network operator, are connected for the first time.
- A user device is switched on and searches for networks in its vicinity.
- A PAN needs to cooperate with a remote AN.

3.1.2 Discovery / Advertisement

The Discovery/Advertisement phase serves to find and select other ANs, available for composition. The network can either place an active advertisement, offering certain services and resources to other networks, listen to advertisements by other networks, or actively discover its neighbors.

This phase also serves to determine the basic connectivity parameters. Information is exchanged about availability of network interfaces, authentication methods, encryption methods et cetera [15].

3.1.3 Security and Internetworking Establishment

In this phase the merging networks must identify themselves and establish a secure internetworking connection. This can be done in a variety of ways, for example by using third party certificates. In principle, all conventional security methods may be used here. This phase is carried out by the Security-FA.

3.1.4 Composition Agreement Negotiation

In a network composition, the involved networks must settle on a Composition Agreement. This agreement is negotiated over the ANI.

There are two possible approaches for the relation between the C-FA (Composition Functional Area) and the other FAs in these negotiations: *Centralized* and *Distributed*. At the moment, both approaches are still being considered. The advantages and disadvantages of these approaches will be discussed in section 3.5.

In a centralized negotiation, the C-FA handles all negotiation tasks. It communicates with the other FAs in the ACS about their part in the composition.

In a distributed or *decentralized* negotiation, each FA negotiates independently with its remote peer, about the functionality they are responsible for. The role of the C-FA here is to combine these partial agreements after they have been agreed upon by the individual FAs and form a realizable umbrella CA. This may sometimes take some mediation between the FAs, because their parts will not always be independent from each other.

The negotiation process takes the same, iterative form for both the centralized and distributed approach:

At first, network A sends network B a proposition (in the distributed approach, partial) CA. AN B reviews this proposition and either confirms to AN A that this is the (partial) CA they will use, or sends a new proposition if it does not agree on any of the points in the CA (part). This process is repeated until either a CA is agreed upon, or one of the networks decides to abort.

The above mentioned proposition can be created in two ways: it may be either a predefined standard CA or a CA built up from a template. The final content will then be dictated by the specific FAs and local policies.

3.1.5 Composition Agreement Realization

Finally, the actual composition is conducted. In this phase, each of the ANs reconfigures its own control functions and policies and, in most cases, a new ACS is formed, according to the settled CA. In some cases, addresses may need to be re-assigned and re-organised.

After composition, the composed ANs keep on exchanging some of the information involved in the composition process. This could lead to a re-iteration of the composition process, or even a decomposition. If, for instance, a new service arises, this would be reason for the other AN to re-initiate the composition.

3.2 Decomposition

If, for some reason, one of the ANs inside a composed AN decides it doesn't want to be a part of the composed network anymore, decomposition takes place.

There can be a lot of different reasons for this. One, obvious reason is losing the physical connection, for example when a person carrying a PAN is physically moving away from his car's VAN (Vehicle Area Network). Other reasons can be a change in the shared services of one of the other networks, or the appearing of another AN offering the same services with better QoS or cheaper.

Because the composed networks offer their services as one single AN to the outside, in most cases, the remaining network will need a configuration update.

3.3 Formalization of the Negotiation Process

I have formalized the Composition Agreement negotiation process in the form of Finite State Automata (FSA). I have formalized both the distributed and the centralized form of negotiation, as described in section 3.1.4.

3.3.1 Distributed Approach

Figure 3.2 depicts the formalization of the actions of the Composition Functional Area (C-FA) in a distributed composition.

Next to each arrow is a text consisting of the preconditions to make this transition (left of the slash) and the actions taken at this transition (right of the slash).

The rectangles, entitled *M(FA)*, represent constituent FSA. These all represent instances of the FSA depicted in figure 3.3. This FSA implements the negotiations between the corresponding FAs of the composing networks in a distributed composition. The arrow from the top-right corner of a rectangle means the sub-FSA finished in a final state, which is not the error state; in other words, the negotiations succeeded. The arrow from the bottom-right corner of a rectangle means the sub-FSA finished in the error state.

The & sign means that all the incoming conditions must be satisfied; in other words, all sub-FSA must have finished in a final state, other than the error state. The $\vee$ sign means that at least one of the incoming conditions must be satisfied; in other words, one of the sub-FSAs has finished in an error state.

So, as can be seen in the FSA in figure 3.2, when an AN is available for composition, first a composition request is sent. If this request is declined or a timeout occurs the composition is aborted. If the request is accepted all FAs are triggered to start negotiating with the corresponding FAs of the AN we are composing with.

This is formalized in the FSA in figure 3.3. Once the FA is triggered by the C-FA to start negotiating, a CA is built and sent to the other party. The negotiation round counter is set to 1. If this CA is accepted then we can go directly to the goal state. If it is not, another negotiation round is needed. If the proposed CA is not accepted by the other party, it sends us a new proposal, so we will have to examine that. If the other party's proposal is acceptable to us, we notify them and finish. If its not, we make a new proposition and send it back. This continues until either an acceptable CA is formed, a timeout occurs or the negotiation round counter exceeds its maximum.

If one or more of the FAs could not agree with the corresponding FA of the network we are composing with, we abort. If all FAs have agreed on their part of the CA, the C-FA needs to check if these parts are independent of each other and, if not, are consistent. Sometimes some mediation may be necessary, so in those cases the involved FAs are told about the problems in their contributions and the negotiations for these FAs start over. If the CA is consistent, we authorize it, send it back to the other party and wait for their authorization. Once we receive this, the negotiations are complete and a CA has been successfully agreed upon, so can now be implemented.

The AN we are composing with can at any time abort the composition, which means these FSA are incomplete. In fact, there should be an arrow from every state in the diagrams to the error state to catch this exception. Since this would largely decrease the readability of the diagrams I have left this out.

Figure 3.2: Formalization of the actions of the C-FA in a distributed composition

Figure 3.3: Formalization of the actions of individual FAs in a distributed composition

3.3.2 Centralized Approach

In figure 3.4, a formalization is given of the actions of the C-FA in a centralized composition.

In the centralized approach, if an Ambient Network is available for composition, first, a composition request is sent. If the request is declined or a timeout occurs, we abort. If the request is accepted the C-FA triggers all other FAs to generate their part of the CA. Out of these parts, a CA is formed and sent to the AN to compose with. Negotiation rounds, similar to those in the distributed approach now start, with the difference that the C-FA now handles all negotiation, instead of each FA apart from the others. Once a CA is agreed upon it is authorized by both parties and the composition can now be implemented.

As in the previous formalization, the AN we are composing with can at any time abort the composition, which means these FSA are incomplete. In fact, there should be an arrow from every state in the diagrams to the error state to catch this exception. Since this would largely decrease the readability of the diagram I have left this out.

Figure 3.4: Formalization of the actions of the C-FA in a centralized composition

3.4 Prototyping Composition

In an attempt to gain some more insight in the principle of composition and to test its feasibility I have built a prototype of a composing Ambient Network. The emphasis in this prototype will be on the CA negotiation process, although it will be set up as two complete merging ACSs. It will be written in Java, version 1.4.2.

3.4.1 Prototype Design

The prototype will be a program in which the user will be prompted to input the actions of his choice. He will be able to create FAs and set-up their services and create ACSs from this. He also will be able to have two ACSs compose and decompose. This will show if this basic composition principle is feasible.

3.4.2 Class Definitions

Six Java classes will be defined to implement this basic composition prototype:

- **Service**
  This class represents a service, provided by a Functional Area, for example 'Internet Access'. This service has a name and a cost. This cost can also be 'infinite', meaning that this service may not be shared. A service can comprise the actual service, or a link to a service from one of the FAs of a constituent FA.
- **FA**
  The FA class represents a Functional Area. It has a name (e.g. 'Connectivity') and it offers a number of Services.
- **CFA**
  The CFA class represents the Composition Functional Area. This is an extended FA class, which includes some actual composition functionality. This class is supposed to handle the compositions as we are prototyping here.
- **CA**
  This class represents the Composition Agreement. Although in practise the CA will be some kind of document, this is represented here by a Java class for easy readability. In this class the details of a composition are stored. It has a common area in which the new AN's name will be stored and a section for each FA where it is stated which services these FAs provide. In some cases a new FA will be built for the new composed network, in other cases an FA of one of the composing networks will be used, especially when one of the ANs does not have this particular FA. This is also stored in the CA.

- **ACS**
  The ACS class represents an Ambient Control Space. It has a name, a number of FAs and a number of constituent ACSs.

- **Prototype**
  The Prototype class contains the main program, which reads and executes the users input.

3.4.3 User Commands

The user will have a couple of commands with which he can build and modulate ACSs.

Commands which can be used to create and modify ACSs:

- **NEWACS <name>**
  This will build a new ACS with the given name. ACS names must be unique, if an ACS with the given name already exists an error will be issued and no ACS will be built.

- **ADDFA <acsname> <faname>**
  This will add the FA with name *faname* to the ACS with name *acsname*.

- **REMOVEFA <acsname> <faname>**
  This will remove the FA with name *faname* from the ACS with name *acsname*.

- **COMPOSE <acs1> <acs2> <comptype> <newname>**
  This will get *acs1* to try and compose with *acs2*. If successful, the name of the new composed ACS will be printed. On failure, the reason for this failure will be printed. *Comptype* determines the composition type. It can be either *centralized* or *distributed*. This determines whether the C-FA handles the CA negotiations, or if it triggers each FA to start its own negotiations. *Newname* is the name of the new composed ACS.

- **DECOMPOSE <acsname>**
  This will decompose the ACS with the given name.

- **PRINTACS <acsname>**
  This command will print all data for the ACS with the given name.

Commands which can be used to create and modify FAs:

- **NEWFA <name>**
  This will build a new FA with the given name. FA names must be unique, if a FA with the given name already exists an error will be issued and no FA will be built.

- **ADDSERVICE <faname> <servicename> <cost>**
  This will add a service by the name of *servicename* to the FA with name *faname*. If a service of this name already exists it will be overwritten, so this function can also be used to change a service. Cost is either an integer, representing the cost in cents per minute for this service, or 'infinite', meaning this service may not be shared.
- **REMOVESERVICE <faname> <servicename>**
  Remove the service with name *servicename* from the FA with name *faname*.
- **PRINTFA <name>**
  This command will print all data for the FA with the given name.

Other commands:

- **EXIT**
  Exits the prototype.
- **?**
  Displays a list of possible commands.
- **LIST**
  Lists all ACSs and FAs.
- **BUILDTESTDATA**
  Command used for quick and easy testing purposes. It builds two ACSs with some FAs, each supplying several services.

3.4.4 Result

In this prototype I have built both negotiation methods and proved them both feasible. Of course this is a very basic prototype and the used methods for accepting and updating (partial) Composition Agreements is very different from the methods that will actually be used, but this prototype was built to test the basic principle of negotiation and has proved this principle to be a good way to achieve an agreement.

3.5 Centralized versus Distributed Negotiation

In this section I will criticise the methods used in the centralized and distributed negotiation forms and give arguments for and against both of them. Both of these options are still open to the Ambient Networks project.

An argument for the centralized form is the fact that the C-FA keeps the entire negotiation process together and remains one hundred percent in control. No decision can be made by any entity other than the C-FA.

This also means that no inconsistencies can arise in the CA, in the contrary to the distributed approach, where in some cases negotiations have to start over for some FAs, because the different CA are not independent of each other, which may take a lot of time.

As in about any distributed method in computer science, the main argument for distributed negotiations is efficiency. Since all FAs can negotiate their parts of the agreement at the same time, they will all finish a lot quicker than a centralized negotiation. Another reason for this can be found in the fact that for every change in the partial CAs, now only a part of the CA needs to be sent to the other party, instead of the entire CA, as in the centralized form. This is a lot more efficient, especially in cases where negotiations will concentrate around one or two points in the agreement, which will probably be a frequently occurring effect.

A negative side of distributed negotiation would be that every FA now also needs to implement a negotiation function. However, this function can probably 'borrow' a lot of functionality from the C-FA, so this is only a minor consideration.

So, in my opinion, at the moment, the distributed form has the preference. However, this depends heavily on the future developments in Ambient Networking. I think there should be a lot of testing in this area to see exactly which of the two options is the fastest in different settings.

It seems that the distributed version would defeat the centralized form in efficiency, because it distributes the negotiation process. However, if a lot of extra iterations of this process have to be made because of inconsistencies in the produced CA, this would seriously slow the negotiation process down.

Chapter 4

Conclusions

In this paper I have summarized the work the Ambient Networks project group has done until now. I have explained about the workings and advantages of Ambient Networks and the concept of composition. I have formalized the process of negotiation that produces the Composition Agreement to give more insight into this composition principle and built a prototype to show its feasibility. I have also discussed the two options for the implementation of these negotiations: centralized and distributed.

Although a lot of the concepts and techniques in this paper may seem a bit as if they came out of a science fiction movie, we are actually quite far along the way of actually living among ubiquitous computers. Also, as I have already mentioned in the introduction, progress in this area is extremely rapid. So, Ubiquitous Computing and Ambient Networking might actually be something that is a big part of our everyday lives in a year or ten, and if everything follows plan, we will not even be aware of it!

Bibliography

[1] G.E. Moore, "Cramming More Components Onto Integrated Circuits", http://download.intel.com/research/silicon/moorespaper.pdf"

[2] IST Advisory Group, "Ambient Intelligence: from vision to reality", ftp://ftp.cordis.lu/pub/ist/docs/istag-ist2003_draft_consolidated_report.pdf

[3] K. Ducatel, M. Bogdanowicz, F. Scapolo, J. Leijten and J-C. Burgelman, "Scenarios for Ambient Intelligence in 2010", ftp://ftp.cordis.lu/pub/ist/docs/istagscenarios2010.pdf

[4] H. Corporaal, I.G.M.M. Niemegeers, F.W. Vaandrager and J.L. van Meerbergen, "Progress White Papers 2006", ISBN-10: 90-73461-00-6, ISBN-13: 978-90-73461-00-0

[5] N. Niebert, A. Schieder, H. Abrmowicz, G. Malmgren, J. Sachs, U. Horn, C. Prehofer, and H. Karl, "Ambient Networks - An Architecture for Communication Networks Beyond 3G", http://www.ambient-networks.org/docs/An_Architecture_for_Communication_Networks_Beyond_3G.pdf

[6] R. Campos et al., "Scenarios for Network Composition in Ambient Networks: a new paradigm for Internetworking", http://www.ambient-networks.org/docs/Scenarios_for_Network_Composition_in_Ambient_Networks_a_new_paradigm_for_Internetworking.pdf

[7] N. Niebert, H. Flinck, R. Hancock, H. Karl, and C. Prehofer, "Ambient Networks - Research for Communication Networks Beyond 3G", http://www.ambient-networks.org/docs/ Ambient_Networks_Research_for_Communication_Networks_Beyond_3G.pdf

[8] L. Fan, N. Akhtar, K.A. Chew, K. Moessner and R. Tafazolli, "Network Composition: A Step towards Pervasive Computing", http://www.ambient-networks.org/docs/Network_Composition_A_Step_towards_Pervasive_Computing.pdf

[9] N. Niebert, C. Prehofer, R. Hancock, T. Norp and J. Nielsen, "Ambient Networks - A New Concept for Mobile Networking", http://www.ambient-networks.org/docs/Ambient_Networks_A_New_Concept_for_Mobile_Networking.pdf

[10] A. Schieder, L. Eggert, N. Papadogluo and F. Pittmann, "Components and Concepts of the Ambient Networks Architecture", http://www.ambient-networks.org/docs/Components and Concepts_of_the_Ambient_Networks_Architecture.pdf

[11] R. Campos, C. Pinho, M. Ricardo, J. Ruela, P. Poyhonen and C. Kappler, "Dynamic and Automatic Interworking between Personal Area Networks using Composition", http://www.ambient-networks.org/publications/ Dynamic_and_Automatic_Interworking_between_Personal_Area_Networks.pdf

[12] D. Zhou, P. Poyhonen, N. Akhtar and C. Pinho, "Ambient Network Interfaces and Network Composition", http://www.ambient-networks.org/publications/ Ambient_Network_Interfaces_and_Network_Composition.pdf

[13] C. Prehofer, N. Papadoglou and M. Johnsson, "On the Evolution of a Naming and Addressing Architecture in View of Dynamic Network Interconnection", http://www.ambient-networks.org/publications/ Evolution_of_Naming_and_Addressing_Architecture_in_View_of Dynamic_Network_Interconnection.pdf

[14] M. Johnsson et al., "Network Composition", http://www.ambient-networks.org/publications/Network_Composition.pdf

[15] J. Colas et al., "Connecting Ambient Networks - Architecture and Protocol Design", http://www.ambient-networks.org/publications/D_3_2.pdf

[16] N. Niebert, "Ambient Networks - A Framework for Mobile Network Cooperation", http://delivery.acm.org/10.1145/1090000/1080779/p2-niebert.pdf

[17] P. Kerssch, R. Szabo, Z.L. Kis, M. Erdei and B. Kovacs, "Self Organizing Ambient Control Space: An Ambient Network architecture for dynamic network interconnection", http://delivery.acm.org/10.1145/1090000/1080782/p17-kersch.pdf

[18] A.S. Tanenbaum, "Computer Networks - Fourth edition", ISBN 0-13-038488-7