

Management Challenges of Context-Aware Services in Ubiquitous Environments

Heinz-Gerd Hegering, Axel Küpper,
Claudia Linnhoff-Popien and Helmut Reiser

Munich Network Management Team, University of Munich, Dept. of Informatics,
Oettingenstr. 67, D-80538 Munich, Germany
{hegering|kuepper|linnhoff|reiser}@informatik.uni-muenchen.de,
<http://wwwnmteam.informatik.uni-muenchen.de>

Abstract. *Ubiquitous environments facilitate the collection of information pieces from sensors, databases, or mobile devices in order to compose the context of entities like users, places, or things. The context obtained in this way can be used to automatically adapt the behavior of services, which results in the new paradigm of context-aware services (CASs). In recent years, a lot of research has covered the functional aspects of CASs. However, CASs in ubiquitous environments impose new management challenges, which has not been considered so far. The goal of this paper is to identify new challenges on CAS management and thus to provide a roadmap for further research in this area.*

Keywords:

Context Aware Service, Quality of Context, Federative Organization Model, Context Information Model, Context Value Chain

1 Introduction

Mark Weiser's vision of *ubiquitous computing* of the year 1991 [18] gains more and more momentum as miniaturization and integration of computing and wireless communication facilities evolve. The Internet has been the enabling infrastructure behind services for individualized information retrieval and new forms of interactivity during the last decade. Now ubiquitous computing is the driving force behind the new service paradigm *context-awareness*.

According to Dey [5] "*context* is any information that can be used to characterize the situation of an *entity*: an entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves". A service then becomes a *context-aware service* (CAS) if its behavior or the content it processes is adapted to the context of one or several entities in a transparent way. This adaptation process will be called *contextualization* in the following. However, before a CAS can be contextualized the context of relevant entities needs first to be obtained from various context sources like sensors or databases, which is called *context procurement*. In recent years, several platforms have been developed for realizing context procurement and contextualization, for example Parctab [14], CAPEUS [13], the ContextToolkit [5], and the Technology for Enabling Awareness (TEA) [15], to name only a few of them.

Many research projects have concentrated on functionality aspects of CAS and mostly regard device-centric applications. However, ubiquitous computing implies the interconnection of many local environments and devices and thus makes it possible that users share their context (and that of other entities) among each other. As a consequence, several actors like users, network operators, service and content providers are involved in context procurement, contextualization, and CAS provisioning. Because ubiquitous computing comes along with highly heterogeneous and distributed environments, our idea is that of an integrated approach of CAS management. Real-time requirements, high dynamic and automation are typical for CASs and evolve new management challenges. But management issues have been neglected or even left unconsidered in former platforms for context-awareness. Therefore, this paper identifies new challenges on a CAS management and thus to provide a contribution for further work in this area.

The remainder of this paper is structured as follows: section 2 introduces an application scenario that outlines the complexity of CASs and serves as a reference example for subsequent sections. To identify management challenges and structure the complex processes of context procurement and contextualization we adopt traditional management concepts[8]. Accordingly, we propose information and organization models dedicated to the special aspects of context awareness and, based on these models, derive the challenges within the functional areas of management. Section 3 covers the information model and introduces a process-oriented value chain to describe context procurement and contextualization in a structured way. The necessity of extending information models with the value chain and context description will be shown. Section 4 covers the organizational model and proposes a role model that identifies the different roles the actors of a CAS infrastructure may adopt according to their functional tasks. This role model is used to identify manager relationships across organizational boundaries. Based on the value chain and the role model, section 5 derives management challenges according to the known functional model, i.e., from the point of view of fault, configuration, accounting, performance, and security management. Finally, section 6 concludes the paper and presents further work.

2 Application Scenario

To outline the complexity of CASs in ubiquitous environments, but also to provide a reference example for discussing management challenges, we introduce the application scenario *Medical Advice and Emergency System* (MAES). MAES is intended for persons with critical diseases or medical disabilities (*patients*). The system gives medical advice to these patients and, in case of emergency, supports the workflow of the rescue crew. Different users of MAES, like patients, physicians, or the ambulance staff, use either conventional mobile consumer devices or special-purpose devices. They are connected to the system via heterogeneous wireless and wired infrastructures. Thus, from this point of view MAES can be seen as a conventional IT service which is implemented by composing interdependent and layered sub-services of different actors like network operators, service and content providers. However, usage of medical IT services like MAES impose very high requirements such as promptness, simplicity, and reliability on the underlying service infrastructure. It is therefore obvious to make these services context-aware in order to reduce awkward interactions between the users and

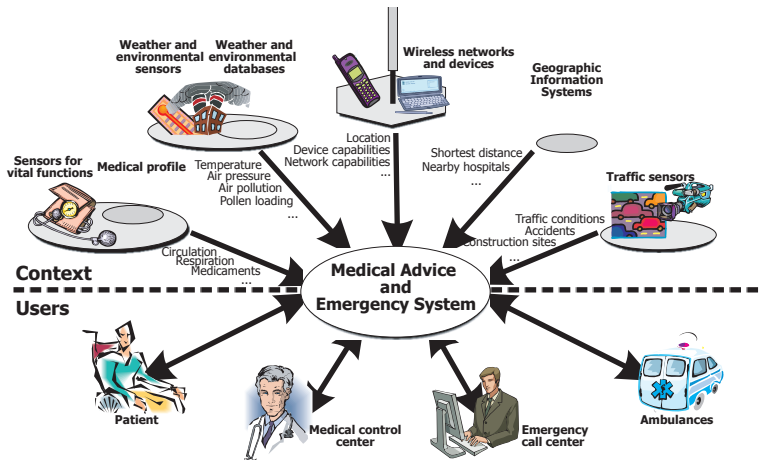


Fig. 1. Users and context information in MAES

the system, to automate workflows, and to adapt these workflows as well as the service's behavior and appearance towards the users according to current circumstances. To achieve this goal, MAES incorporates an extensive set of heterogeneous context information which cover medical, spatial, environmental, and technical aspects. Figure 1 gives an overview of the different types of users and the context information processed by MAES. An important input for MAES is the patient's medical situation, which comprises her current vital functions and her medical profile. Vital functions are derived by sensors which are worn by patients [9] and which deliver information about consciousness, respiration, circulation, and metabolism. The medical profile contains diagnostic findings made by physicians as well as prescribed medicaments and therapies. Depending on the current patient's vital functions and under consideration of her medical profile MAES may trigger the medical advice function or the emergency handling function.

The medical advice function notifies the patient or a remote medical control center about the current medical state and proposes remedial actions, for example to take certain medicaments or to consult a physician. These recommendations are combined with a list of nearby pharmacies and physicians, thus it is necessary to derive the current location of the patient. Weather conditions like temperature, precipitation, and atmospheric pressure, as well as environmental influences like air pollution, pollen loading, and ozone concentration may be considered in addition. Based on this environmental information, MAES can recommend to avoid exhausting activities.

Emergency handling is a workflow starting with the receipt of an emergency call and terminating with hospitalizing the patient. Upon receipt of an emergency call, the operator receives the patient's vital functions and medical history and can select an appropriate hospital with free capacities and in close proximity to the patient. Afterwards, MAES notifies an ambulance which is equipped with the required personal and material configuration (derived from the patient's medical profile) and which can reach the patient as fast as possible. The latter requires to derive the shortest path between ambulance and patient and to consider current traffic conditions along that path. The ambulance is

supplied with the patient's medical profile and, if the patient has not been reached yet, with her latest vital functions.

From this application scenario several features of CASs can be derived that impose high requirements on management tasks. The following list itemizes the most important of these features:

- CASs obtain data from a very heterogeneous set of *context sources*, e.g., sensors, mobile devices, and databases, which are highly distributed in the ubiquitous environment and which are located in the domains of different actors like operators of sensor networks or content providers.
- Data delivered by different context sources is consequently also very heterogeneous and usually varies in its update frequency (if any), its accuracy, and its format of representation.
- Context information is not only the output of some context sources, but the result of service and user specific context procurement, i.e., the processing and distribution of data delivered by context sources.
- CASs are characterized by a higher level of mobility. Not only users are mobile but even parts of the technical infrastructure, e.g., context sources. A CAS should even work if its mobile users leave the service area of her original provider or the coverage area of a certain context source. There are roaming challenges not only for users and services, but also for context information delivered by different actors.
- A context information may have a spatial and temporal validity. Due to the mobility of users, it might be necessary to locate relevant context sources during service usage and to identify resources needed for context procurement in an ad-hoc manner.

Provisioning of CASs will therefore be a new challenge for involved actors and their management systems. CASs will only be accepted at the market if they are easy to use, easy to configure and if the organizational and management boundaries are absolutely transparent from the customers' point of view. This means that realization and management of CASs require a completely new kind of interworking and cooperation between independent actors.

3 Context Information Model

Context-awareness implies the procurement of context and the contextualization of CAS components according to this context. Recent approaches have been dealing with these issues, for example, [5], [15], and [14], slightly differing from each other in the proposed structuring and terminology of these processes.

To automate the context procurement and contextualization we introduce the concept of a *value chain* which covers all steps happening between sensors and CAS components. Classical information models must be extended with a formal description of context and they must represent the value chain to make context procurement and contextualization manageable. We call these extensions context information model, which will be presented in the following.

Figure 2 shows an example of such a value chain which describes the procurement of medical profiles, vital functions, location and weather information as required for the

contextualization of the medical advice function of MAES. Clearly, the universal value chain approach applies to other CAS scenarios, too.

The context of an entity may be derived from very different sources, e.g., sensors checking the patient’s vital functions, positioning methods like GPS for obtaining her current location, or databases containing her medical profile. Capturing data from these context sources is a process we call *sensing*. However, sensing merely provides the “raw material” of context, which is referred to as *low-level context information* in [5], and which is often not interpretable by the requesting CAS. Usually, one or several steps of *refinement* have to be performed in order to derive *high-level context information* as required by the respective CAS. Among other things, refinement comprises the transformation between different formats of representation (e.g., from GPS coordinates to street names and numbers), the extension of context information with attributes (e.g., to express its accuracy), or the combination of context information to derive another one (e.g., calculating the distance between patient and ambulance by using their locations). After refinement, all the required context information which is related to a particular entity and which is of relevance for a particular CAS needs to be allocated, a process which is called *aggregation*. Figure 2 shows possible context information of a patient required for the medical advice function. Finally, the aggregated context is used to contextualize the CAS, which represents the last step of the value chain.

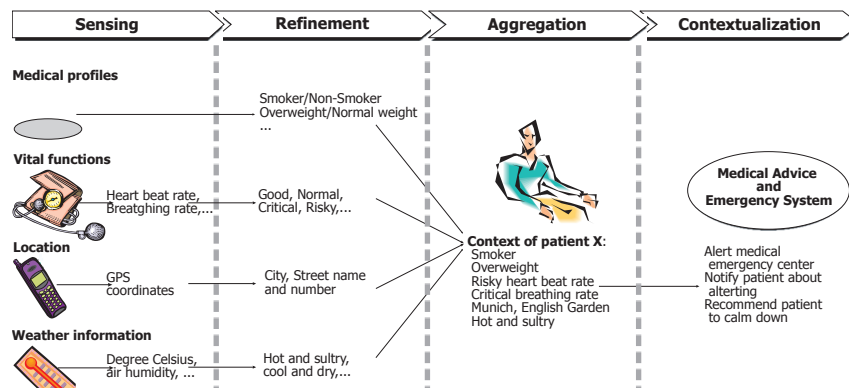


Fig. 2. The Context Value Chain

It must be stressed that a value chain may comprise several sequences of sensing, refinement, aggregation, and contextualization, one for each context information to be considered. These sequences might be executed in a prescribed order or in parallel, either entirely independent from each other or with need for synchronization. The exact coordination of the various sequences depends on a lot of circumstances like the requirements of the CAS, the range of context information to be processed, interdependencies between context information, the availability of context sources, etc. Each step of a sequence may be executed in an iterative manner and may comprise several sub-steps. Furthermore, the steps of a sequence may be triggered on demand or they may be event-triggered, i.e., if the value of a context information exceeds a pre-defined

threshold. Thus, the description given here is, due to simplicity, rather a coarse-grained representation of the value-chain concept.

Building a context chain and selecting the most appropriate context sources is a federative process. Automation of contextualization and using context information across organizational boundaries can only be done if all participants have the same notion of how to interpret the data. Therefore the management information model has to be extended with a *Context Description Language (CDL)* to formalize the specification of context. Without such a formal and automatically evaluable description of provided context information selection of context sources and brokerage of such information is infeasible. A starting strategy on the way to a uniform notion of context is to build context categories, i.e., *device-specific*, *environment-specific* and *user-specific context* [2]. However existing models within these categories are focused mostly on the considered application scenario, e.g., [19] concerns human computer interaction and models context within a hierarchical model which is optimized for database storage with quite few context-classes. A quite narrow and static model for device specific context for resource detection in ad-hoc networks is presented in [12]. An all-embracing model for all context categories is still missing, but nevertheless needed. Therefore a comprehensive CDL had to be specified and information models like the Common Information Model (CIM) should be enhanced with it. Even if an unambiguous context description exists, it will be not fully sufficient for selecting context sources. A formal notion of *Quality of Context (QoC)* as an integral part of each context description is needed to build an appropriate value chain in terms of quality and price of their aggregated context. Examples for such QoC parameters are accuracy, availability, timeliness, validity period, system of units, conversion factors, etc. [3].

As there are high timing constraints in CAS the process of context procurement might not contain any procedure which requires human interaction. Automation can only be achieved if all sub-processes of the value chain can be automated. One of our ideas to investigate in the future will be the modeling of each sub-process of the value chain as a managed object. The management system has to regard the formal context description with actual QoC as input parameters to select the appropriate sub-process(es) with its “best” context source.

4 Federative Organization Model

A nontrivial CAS can only be realized in an interorganizational manner. As seen in section 2 the MAES provider can only offer its service by composing sub-services of various other actors (e.g., operators of sensor networks, mobile network operators, information system providers, traffic management systems, ambulances, hospitals, etc.). Thus, several actors must establish a federation in order to establish the value chain presented in the last section by composing their sub-services. To classify the actors according to their roles they play in such a federation, and thus to structure management tasks from an organizational point of view, we propose a role model for CAS.

In this model, an *actor* denotes an individual, organization, department, or enterprise, which offers services to other actors, which consumes services from other actors, or which does both of them. From a system’s point of view, each actor autonomously operates and controls its own administrative technical domain, consisting for example,

of a network infrastructure, a server farm, or only a single mobile device. An actor may adopt one or several roles. A *role* represents a certain field of activity of an actor and is associated with a certain set of sub-service components for realizing and controlling the value chain. Figure 3 shows our approach of a role model for a CAS infrastructure. Note that our understanding of a role model is in accordance with that of business models that have been proposed for telecommunication systems like TINA [17], UMTS [1], or the MNM service model [7]. However, we prefer to define a dedicated role model in order to be independent of a particular network technology and to highlight the special problems and tasks of context-awareness.

The central role of our model is the *CAS provider*, which creates and deploys CASs like MAES and offers and sells them to a *CAS customer*. The CAS customer interacts with the CAS provider in order to negotiate service level agreements and to control customer service management on behalf of one or several *CAS users*.

The CAS provider obtains context information for contextualization of its services from a *context provider*, which is usually the operator of context sources. For example, a context provider may be the operator of a cellular network, which tracks a CAS user and delivers that user's current location to the CAS provider. In some cases more rules can be adopted by the same actor, e.g., the roles of context provider and CAS user are adopted by the same actor. For example, a patient uses the MAES and simultaneously delivers her vital functions to the system.

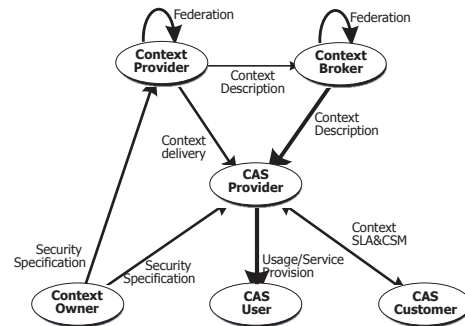


Fig. 3. The Role Model of a Context-Aware System

Due to the heterogeneity and diversity of context information, it is unlikely that a certain context provider is able to deliver all context information required for a CAS. Also, a context provider may only cover a limited geographic area, and due to the mobility of CAS users the relevant context providers are not known in advance, but must be identified during CAS usage. The *context broker* supports this identification. It maintains a directory for registering context descriptions of information a context provider is able to deliver together with its QoC. This directory can then be requested by CAS providers in order to find appropriate context providers.

For many services it would be desirable that a CAS user has access to context information which is related to another actor. For example, in the MAES scenario the emergency control center requires the patient's vital functions and medical profile in order to decide about the right ambulance configuration. From a security point of view, this a very sensitive matter, because an actor must always have control about the processing and accessing of her context by other actors. It is therefore inevitable to establish the role of a *context owner*, which represents an entity context is related to and which is able to specify access restrictions regarding its context.

The value chain presented in the previous section is realized by the interactions between different context providers on the one hand as well as between context providers and a CAS provider on the other. The exact mapping of sensing, refinement, aggregation, and

contextualization onto the roles is a point for further discussion. However, it is obvious at least that context providers always performs the sensing, whereas CAS providers are responsible for the contextualization. Whereas the value chain is directly realized by the roles of context and CAS providers, the remaining roles have nevertheless a significant impact on the value chain, which is to be discussed within the scope of the management functional areas in section 5.

As can be derived from this role model, each actor will have its own technical infrastructure and management system(s), and there is no "leader" which can intervene regulatively. Enterprise management concepts with the aim to build an integrative management umbrella on top of various different management systems is infeasible, because in such a federative environment there will be no actor prominent enough to establish and operate such an umbrella system. Most of the participating actors will not allow deep influence within their own sense-making process. Therefore, managing CAS will only work in a federative and not in any kind of hierarchical organization model, i.e., operation and management of a CAS infrastructure will only be achieved in an inter-organizational cooperation between the various actors. The managers of different actors need to interact among each other in order to establish and control the value chain according to the management functional areas. However, in such a federation each actor may deploy its own management system, including actor-specific representations of resources and sub-services (managed objects) as well as management protocols. Thus, management gateways are required at the boundaries of each actor's technical domain in order to mediate between the heterogeneous management systems.

5 Management functional areas

In contrast to conventional, non-context-aware IT services the realization of management functions including the inter-organizational mediation process is a hard and troublesome matter. The mobility of users in combination with the availability of appropriate context sources forces a high degree of flexibility and automation on management functions during CAS provisioning. In addition, these management functions may be subject to hard real-time requirements. The impact of these requirements is discussed in the subsequent sections for each of the functional management areas separately.

Configuration The concept of context value chains raises the question how to establish them if a certain context information is needed. Obviously, its configuration depends on the context information itself, whether or not it is based on other context information, and the requirements of the CAS. Whereas there are some static aspects, configuration is frequently based on dynamic criteria, above all the mobility of users. For example, the availability of context sources depends on the user's location, i.e., if the user moves, the set of available context sources may change accordingly. Furthermore, a context information can usually be derived by different value chains, differing from each other. Establishing the most suitable value chain for each individual user — taking into account her current situation — will be the main configuration management challenge for the CAS provider. For this purpose the CAS provider needs an integrated configuration mechanism across heterogeneous systems and across various context providers. This means that each context provider has to make available an automated configuration interface for its CAS providers. Beneath building the value chain its automatic adaptation to a constantly changing infrastructure, e.g., if sensors vanish, will be another

challenge. The CAS provider itself must support its customers with a customer service management (CSM) interface [10]. A CAS must be easy to use and easy to configure. The configuration simplicity will be the most important aspect for CSM.

The service lifecycle as presented in [7] is feasible for establishing a CAS infrastructure in the long term. However it does not cope with the ad-hoc creation of an user individual peculiar instance of a certain CAS regarding the current context of the context owner. Therefore during the usage phase of each CAS a micro-lifecycle has to be established and traversed for building each individualized service instance. This micro-lifecycle contains all phases of the long term life cycle, i.e., design (resp., adaptation), negotiation, provisioning, usage and de-installation (resp., de-allocation). Constitution and cycling the micro life cycle must be supported by configuration management functions.

Fault The challenges in fault management cover two areas: faults in sensing and faults in the procurement of context. Context sensors are often proprietary highly specialized devices with their own raw data format, limited usage interfaces, limited CPU as well as limited memory, and without management or alarming functions. Fault detection and recovery in sensor networks therefore is not trivial. Management mechanisms have to cope with e.g., power management, radio energy management, sensor location, (collaborative) signal processing, sensor synchronization, etc. [6]. Context sensing within CAS is characterized by highly dynamic and short-term update cycles of context information. Especially, for sensor networks self-healing and self management mechanisms are needed to be able to cope with dynamics and the technical deficiencies of sensor networks — beneath their pure sensing functionality — must be compensated by the management system.

Sensing — capturing data from context sensors — does not always work; this is a technique intrinsic fact within context sensing. A simple example for that is the impossibility of GPS localization without a direct line of sight to some satellites. For localization in CAS this means that a GPS sensor is useless indoor. The effect of such technical restrictions for the process of building a value chain is the same as losing a certain sensor. The fault management system must be able to detect this (temporary) absence of sensors and must support the configuration management as well as activate the micro-lifecycle to build an alternative value chain with nearly the same QoC.

Accounting The accounting management of CAS must be able to deal with the fact that each user gets its own and unique service which will be individually customized for her. In the worst case such a service will only be used once. Developing of fair tariffs for a user individual CAS which will be used only a few times is an unsolved problem. From the customer's point of view the predictability of tariffs and keeping the costs of CAS under control is mandatory.

The accountable units, the mechanisms, the measurement points for gaining them and their exchange within an interorganizational CAS infrastructure have to be defined. Both sides — customer and CAS provider — need a concept for conservation of evidence for the CAS usage and accordant accountable units. The CAS provider needs new cost and charging models for context information and for CASs. These models must then be deployed between different (sub-)providers which contribute information to the context value chain.

Concepts of calculating costs, negotiating them with the customer, and concluding them in an agreement before actually using the CAS will not work in such highly dynamic

environments. For CAS ad-hoc agreements and ad-hoc pricing are necessary. It can not be assumed that each customer will apply and subscribe before actually using a service. Rather, a customer would like to define a price ceiling for a certain CAS and awaits best possible service delivery within her price range. From management point of view this connotes a complete inversion: not the service defines the price but the price causes the type of service provisioning. QoS based pricing concepts can be a starting point (for a review cf., [16]), however concepts for a price-driven service adaptation, building the value chain and price-driven service-provisioning are fundamentally needed.

Besides the classical roaming in cellular phone networks CAS users will claim for "service roaming" which enables a seamless CAS usage even if the user leaves the service area of its CAS provider. For that purpose and to follow the principle "one face to the customer" an inter-provider accounting system is necessary.

Performance A CAS "performs well" if its user will always get appropriate QoS. A CAS provider must be able to build and even update value chains with appropriate QoC. In provisioning of CAS, in procurement and update of context information real time requirements can not be eliminated. For building a certain context value chain this implies concepts for optimizing this process in terms of response times. Therefore notions for optimal context caching and context reuse can be helpful. Reuse of sensor data, context or of a preliminary stage within a value chain, can shorten allocation process. A basic QoC-attribute therefore will be the validity period of each context information. Without this caching is not feasible. Because of the frequently changing circumstances during CAS usage and as changes in the actual environment of a context owner might change its context a predefinition of a validity period is not that simple.

Security Providing context aware services requires foremost the same security mechanisms — e.g. for identification, authentication, and authorization of participating authorities, non repudiation, etc. — as traditional services. However there are some new aspects respectively tighten requirements especially regarding confidentiality, integrity, availability (CIA), informational self-determination, privacy, and trust level management.

Context information and according profiles might be highly sensitive and private information for each user; e.g., medical profile or current medical status in the MAES scenario. It may be disastrous for the context owner if even parts of this information came to wrong hands, would be manipulated, or if the context sources the service depends on would be unavailable. Existing confidentiality mechanisms do not scale well for CAS, because protect-able data might be generated by a huge amount of context sources under responsibility of lots of different context providers and transmitted via various communication mechanisms. Classical confidentiality mechanisms cope with a manageable group of beforehand known communication partners or they are designed for certain communication channels (e.g. IPSec, ssh, WEP). In CAS none of this prerequisites are entirely fulfilled. There are contractual relations between the context owner and some CAS providers, but normally there is no such contract between her and the various context providers. Even worse, the context owner often does not even know context providers which generates information about her. The information flow is unknown and uncontrollable for her. Data is transmitted via heterogeneous channels (BAN, WLAN, 3G Networks, TCP/IP Networks, etc.). Therefore integrated and inter-organizational security mechanisms are which can cope with and configure different

heterogeneous security mechanisms on different and heterogeneous infrastructures in a consistent manner.

Similar problems arise regarding integrity, accuracy, and availability. The CAS user must be absolutely sure that context information about a context owner delivered by herself, by the CAS provider, or even directly by the context providers are reliable, accurate, and on time. For proper CAS operation not only the availability of servers or network components are critical, but also the availability of context sources. However sensors are not designed to have high availability and it is easy to disturb them physically. Mechanisms for automatically selecting and querying alternative context sources with similar QoC are essential for the CAS provider in case of unavailability. All participating roles must be able to verify the integrity-aspects (reliability, accuracy, timeliness) of received data independently of other roles.

For the context owner the most important security requirements are her informational self-determination, her privacy, and controlling the hazards of personalized context information. Obtaining condensed context information might enable serious misuse. Potential dangers are user tracking and sophisticated user profiling. Concepts for the context owner are necessary to specify and perhaps even control information flows and to determine who can see or use her actual and historical context information. This means that the context owner must be enabled to authorize and prohibit access to data which is not under her own control. For privacy aspects within web-services a privacy policy solution has been presented in [4,11]. Parts of these concepts might be applicable or enhanced for CAS. But specification of such access rules must be done easily and independently of arbitrary access control models used at the provider side. Protectable data is produced by a huge amount of different providers. Even worse for the context owner, there are unknown context providers which are known only by the CAS provider. In such cases, and in order to relief herself, the context owner might delegate the informational self-determination tasks to the CAS provider or to a security provider. Therefore and as a basic security service, a formal trust model and its implementation within a trust level management system could be helpful. Today, trust is mostly defined implicitly and in a "make-or-break" manner. If it is feasible to define, specify, and check the level of trust, then it will be calculable and automatically processable. For services which need the user context without personalized information a concept for anonymous service usage should be available to avoid user profiling.

6 Conclusions and Future Work

In this paper new challenges caused by CAS have been identified and presented. No doubt, more questions and problems are raised than can be answered at the moment. Some approaches to cope with this challenges have been presented. Within CAS scenarios with real time requirements, i.e., with short time and user-individual service adaption needs and with different organizational domains there is an urgent necessity for self-management concepts. In coping with sensor networks and their management deficiencies concepts for self-healing will be helpful.

This work is intended as a starting point for lots of research which must be done in the future. Our next steps in this area will be the observation of further CAS application scenarios to identify the functional building blocks for CAS-specific FCAPS. A

deeper investigation of interorganisational aspects is planned regarding the interactions which are relevant for management purposes. Furthermore we are building managed object classes for the role model, the processes contributing to the value chain and for context information. We are working on the specification of a CDL and as QoC will be a prominent part we try to identify universal QoC parameters for the different context categories.

Acknowledgment

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of the paper. The MNM Team directed by Prof. Dr. Heinz-Gerd Hegering is a group of researchers of the University of Munich, the Munich University of Technology, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. For more information see <http://www.nmteam.informatik.uni-muenchen.de>.

References

1. 3GPP Technical Specification Group Services and System Aspects. Service Aspects; Stage 1 Service Requirements for the Open Service Access (OSA) - Release 4. Technical report, 3rd Generation Partnership Project (3GPP), March 2002.
2. N. Anerousis. Pervasive Computing. Tutorial at the 8th IFIP/IEEE International Symposium on Integrated Network Management (IM 2003), March 2003.
3. T. Buchholz, A. Küpper, and M. Schiffers. Quality of Context Information: What it is and why we need it. In *Proceedings of the 10th HP-OVUA Workshop*, volume 2003, Geneva, Switzerland, July 2003.
4. Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Recommendation REC-P3P-20020416, W3C, April 2002.
5. Anind K. Dey. *Providing Architectural Support for Building Context-Aware Applications*. PhD thesis, Georgia Institute of Technology, November 2000.
6. D. Estrin, A. Sayeed, and M. Srivastava. Wireless Sensor Networks. Tutorial at the Eighth ACM International Conference on Mobile Computing and Networking (MobiCom 2002), September 2002.
7. M. Garschhammer, R. Hauck, H.-G. Hegering, B. Kempter, M. Langer, M. Nerb, I. Radisic, H. Rölle, and H. Schmidt. Towards generic Service Management Concepts — A Service Model Based Approach. In G. Pavlou, N. Anerousis, and A. Liotta, editors, *Proceedings of the 7th International IFIP/IEEE Symposium on Integrated Management (IM 2001)*, pages 719–732, Seattle, Washington, USA, May 2001. IFIP/IEEE, IEEE Publishing.
8. H.-G. Hegering, S. Abeck, and B. Neumair. *Integrated Management of Networked Systems — Concepts, Architectures and their Operational Application*. Morgan Kaufmann Publishers, ISBN 1–55860–571–1, January 1999. 651 p.
9. C. Kasabach, C. Pacione, J. Stivoric, A. Teller, and D. Andre. Why the Upper Arm? Factors Contributing to the Design of an Accurate and Comfortable, Wearable Body Monitor. Technical report, BodyMedia, 2002.
10. M. Langer, S. Loidl, and M. Nerb. Customer Service Management: A More Transparent View To Your Subscribed Services. In A. S. Sethi, editor, *Proceedings of the 9th IFIP/IEEE International Workshop on Distributed Systems: Operations & Management (DSOM 98)*, Newark, DE, USA, October 1998.

11. Marc Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. In G. Borriello and L.E. Holmquist, editors, *4th International Conference on Ubiquitous Computing (UbiComp2002)*, number 2498 in Lecture Notes in Computer Science (LNCS), pages 237–245. Springer, 2002.
12. G.-C. Roman, C. Julien, and A. L. Murphy. A Declarative Approach to Agent-Centered Context Aware Computing in Ad Hoc Wireless Environments. In *The 1st International Workshop on Software Engineering for Large-Scale Multi-Agent Systems (SELMAS'2002)*, May 2002.
13. M. Samulowitz, C. Michaelles, and C. Linnhoff-Popien. Adaptive Interaction for Enabling Pervasive Computing Services. In *2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE 01)*, Santa Barbara, California, USA, May 2001. ACM.
14. B. N. Schilit. *System Architecture for Context-Aware Mobile Computing*. PhD thesis, Columbia University, New York, 1995.
15. A. Schmidt and K. van Laerhoven. How to Build Smart Appliances. *IEEE Personal Communication*, pages 66–71, August 2001.
16. B. Stiller, P. Reichl, and S. Leinen. A practical review of pricing and cost recovery for internet services. In *Netnomics – Economic Research and Electronic Networking*, volume 3. Baltzer, The Netherlands, March 2001.
17. TINA Business Model and Reference Points - Version 4.0. Technical report, TINA Consortium, May 1997.
18. Mark Weiser. The Computer of the 21st Century. *Scientific American*, September 1991.
19. H. Wu, M. Siegel, and S. Ablay. Sensor Fusion for Context Understanding. In *Proceedings of IEEE Instrumentation and Measurement Technology Conference*, Anchorage, AK, USA, May 2002.