Cooperative Mobile Ambient Awareness

Theo Kanter

theo.kanter@ericsson.com, Ericsson Research, SE-16480, Kista, Sweden

Abstract - *In this paper, I examine the requirements, conditions, and work in progress regarding support for cooperative mobile ambient awareness in heterogeneous wireless and mobile infrastructure.*

I. Introduction

Context is an important enabler for context-aware applications implicating new ways for users to interact with mobile services. Context is commonly understood as an abstract concept representing the *user's* situation and intentions. Mobile devices can deduce and synthesize context information out of aggregated data from a number of primitive sources such as sensors (e.g., GPS) or mobility management. Thus (user) context can be used to either present content to the user, which is relevant to the user's situation, or change the mode of interaction in order for the user to be able to focus on relevant content.

A less common but equally important interpretation of Context is *infrastructure* context, which represents user-related parameters concerning what the infrastructure is trying to achieve in order to deliver content to the (mobile) user. Service delivery can thus be made context-aware, directly related to infrastructure context and indirectly related to user context. Infrastructure context is aggregated and computed from a similar set of data as user context but has a lower level of abstraction pertaining to how achieve service delivery.

Both user context and infrastructure context are important enablers for enhancing the performance of mobile network service delivery. Context can be used to create awareness of the communication context of users and leverage this to make more optimal choices for use of available connectivity regarding content sharing between users [1,2,3]. This is especially important in heterogeneous wireless networks with wide area and local area communication paths with different cost profiles and data rates, which may be available or unavailable, which may be hard to predict. However, prediction is possible when we enable mobiles to aggregate, process, and report context data such that they can use this knowledge to make better decisions with respect to use of available communication for sessions or sharing of content.

Equally importantly, by cooperating and sharing context information about observed entities and infrastructure mobile users can prepare for interaction with entities that they are likely to discover *prior* to their arrival in the area that these entities are located in. Hence, we want to enable ambient awareness support, which enables users to focus on goal-directed interactions in certain situational contexts, and thus:

1) *minimize* or eliminate required user attention for managing dynamic changes in the user's computing and communication environment and

2) *maximize* user attention for and interaction with entities/resources/artifacts that are relevant to the user's goals and situational context.

II. Scenario

We envisage a communication space with people, artifacts, and places. People, artifacts, and places can be extended with computing and communication capabilities and thus be present in this communication space as entities that are able to carry, exchange, and extend local knowledge about other entities and knowledge about how to use these entities.

Stationary entities can for instance be so-called hotspots, i.e., local areas with high data rate wireless connectivity, although we could envisage moving networks such as Wireless LAN access points installed in vehicles. Hotspot areas can contain a number of entities or be visited by moving entities, e.g., users. Typically, users will

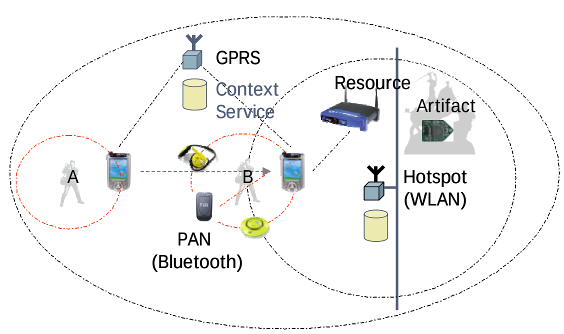

Fig. 1. Example communication scenario

move while most other objects are stationary.

Fig. 1 shows an example communication scenario with users A and B who carry several personal devices that interconnected in a Personal Area Network (PAN). At least one of the users' devices has several other wireless interfaces for local area network connectivity (e.g., WLAN) and wide area network (e.g., GPRS) connectivity. User B is visiting a hotspot that has a common resource (a projector with a wireless interface) and an artifact (e.g., a room) with a multimedia presentation attached to it, achieved via an association between the two in the Hotspot entity. Users A and B are members of a group that is registered in a Context Server. A and B can leverage access to multiple wireless networks for communicating. B is visiting the hotspot and B obtains knowledge about entities in that area (hotspot access, a projector, and room with art).

A is not in that area, but as A and B are members of the same group they can share context information that prepares A for interacting with the entities that B has context knowledge about but A does not. This is important as now A can either communicate with B and share entities that are near B but A has no knowledge of, but it is also important because A can ask fellow members of a group whether they possess context knowledge that is relevant to where A is moving; irrespective of whether A wishes to converse with other group members.

The key idea is that while these are two independent entities, the user can interact with them in such a way that they can use the projector (with a built-in audio output device) to view the multimedia presentation about this painting/sculpture/... both locally and share this information with other members of the group who are non-local.

III. Approach and Requirements

In this section we discuss our approach and requirements, further exploring the conditions and requirements for enabling cooperative mobile ambient awareness.

Cooperative (mobile) behavior requires us to scope interaction between entities as otherwise we would other create a situation where all entities in the whole world would be expected to have or provide knowledge about all other entities, even in areas that were never and will never be visited by the user. Therefore, we require context management to exploit the concepts of *locality* and *social & physical groups*. This has the additional benefit of preventing communication patterns (such as global broadcasting) that do not scale well in wireless networks.

We wish to exchange context information between related entities, which enables them to cooperate as outlined in the scenario. The required *exchange protocol* must support local communication, which also can include peer-to-peer communication, so as to avoid unnecessary signaling via a central server. A context exchange protocol must be generic with respect to the semantics of the exchanged context as these semantics may change over time. This implies that the protocol should provide mechanism(s) for altering and extending its repertoire of context events and context ontologies and should provide support for disseminating context events and context ontology changes to other entities.

Entities carry a data cache containing a context ontology and model of observed entities in the communication environment along with entity profiles and relations between these entities. This model has a dual purpose of being descriptive (enabling the entity to tell other entities what it knows about its environment) but also interpretable (enabling the entity to respond intelligently to context events that it may receive. Our choice to use scoping allows users (and devices) to be aware of knowledge associated with visited areas that can be discarded when not used or used infrequently. Additionally, the user may join or leave groups, hence causing inconsistencies between the data caches of members; therefore, the exchange protocol must support synchronization of knowledge relevant to the group between entities currently in this group.

The protocol should allow (some) stationary entities to voluntarily provide access to their context data caches, thus acting as local context repositories, these repositories can provide this knowledge as a service to other entities (which may in some cases include providing it to those who are not in the same location or same group). Whereas mobile devices usually have quite limited resources (e.g., computing, power, bandwidth,

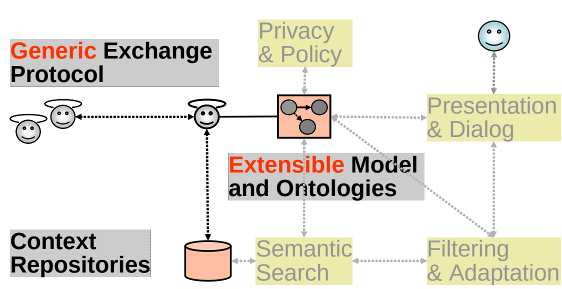

Fig. 2. Logical overview

etc.), stationary devices do not have such limitations and could aggregate unlimited amounts of data. Additional advantages of using context repositories are that they can:

- provide data mining services (that can be made available and browsed by users or service operators)
- offload mobile entities with limited resources

From these above requirements we also derive the need to support privacy and policies, e.g., for sharing information between group members and groups and disseminating context information to other areas. The contributions with respect to creating support for cooperative mobile ambient awareness in comparison to previous work are support for scoping with respect to areas and user groups and more importantly, utilizing these scoping concepts it facilitates cooperation in order for entities to:

- include other entities in the conversation when these other entities are present and known in the local area of one of the conversing entities
- prepare for interaction with entities prior to arrival in a local area where these entities are present.

IV. Context-Aware Middleware

In this section, I briefly discuss the merits of middleware technologies for pervasive computing in relation to the approach. Bluetooth SDP, JINI, Salutation, UPnP, and SLP [4] feature discovery mechanisms. Bluetooth SDP concerns device discovery over Bluetooth links (L2CAP) and is therefore not relevant. JINI and Salutation uses lookup servers and object brokers, instead of the required peer-to-peer model. UPnP uses SSDP which is limited to discovery in a single subnet and cannot be used in large scale mobile networks. SLP uses a directory agent but does not require it, in which case SLP uses multicast for service discovery. However, SLP does not address how end-points should communicate. Therefore, our approach is to introduce similar capabilities for Session Initiation Protocol (SIP) [5] and extend its eventing framework [6] with discovery mechanisms. SIP [5] supports addressing of entities (e.g., users, etc.) and invocation of communication sessions between peers. Theo Kanter's dissertation [7] describes how to extend a SIP User Agent and use it as an Agent that represents entities. The advantage of the approach is that SIP support for locating hosts and invocation of services allows us to move the point of service integration out to the end-points. This can be contrasted to the use of service object brokering such as in [9], which limits possibilities for end-points to negotiate services on a peer-to-peer basis.

V. Early Results

In our pilot project at the Wireless Center, we built a prototype of a system, which implements important features of our approach, discussed below. The prototype (shown in Fig. 3) comprises mobile devices that have multiple wireless interfaces (GPRS, WLAN) for accessing the Internet and (at least) one context server that offers a data store and query interface for context information reported to it by mobile devices. Devices carry an Agent, which represents entities (e.g., users, etc.), a Session Initiation Protocol - User Agent (SIP-UA), and a Context Engine. We enhanced the SIP-UA to deal with new SIP service-discovery extensions [2], which we created to support discovery of entities in the local and wide area. The SIP-UA can act as Presence Agent [8] which has access to context data via a CDXP client. This is particularly important as this implies that also the Agent has access to services. The Context Engine consists of:

- a Context Knowledge Base (KB)
- Sensor Management with pluggable interfaces and wrappers (with data fusion capabilities) for physical and logical sensors
- optional LDAP support for remote access of context information in the Context KB

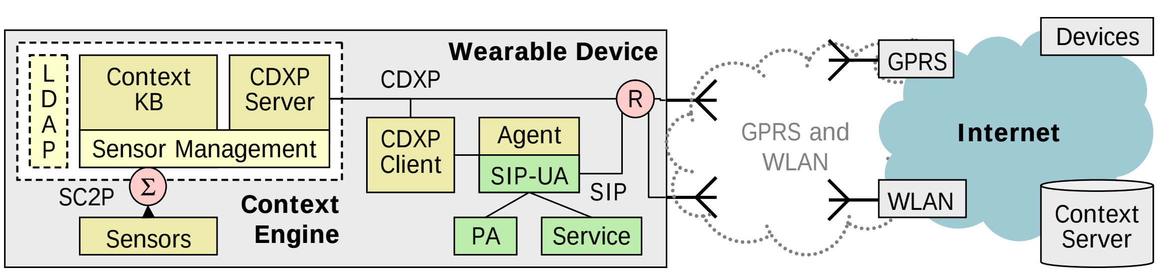

Fig. 3. Prototype Overview

- a novel Context Data eXchange Protocol (CDXP) and Server for the exchange of context data with other entities (both locally on the device or remote).

A. Service Discovery

Service discovery enables users to find services offered by other hosts in the network. The enhanced SIP User Agent (SIP-UA) understands new extensions (DISCOVERY, ACCEPT, and DENY) [2] to the SIP eventing framework [6] in order to register certain services in the context server or find peers that registered services. Peers send discovery messages via unicast and optionally via multicast in a local area network. Messages are expressed in XML, and may reference entities and services that are organized in taxonomy trees (e.g., as shown in Fig. 4).

B. Context Data

Taxonomy trees offer a sufficient model for organizing context information as identified by Held et al. [10]. Taxonomy trees enable unambiguous naming and identification of context information, services, and entities), as well as dynamic binding and extensibility. Context information is identified by the path from the root to the correspondent node. Context Data is aggregated from various sensors via novel sensor sampling protocol SC2P [1] & Sensor Management that has data fusion capabilities, which combined decouple the Context Engine from the sensors (see Fig. 3). Besides physical sensors on the device itself [11], various types of infrastructure information can contribute to context that is offered to other hosts.

C. Context Data Exchange

The Context Data eXchange Protocol (CDXP) [1] is a light-weight XML-based protocol for exchanging context data between peers. CDXP can use both UDP and TCP as transport. CDXP employs a well-known port for the CDXP server. CDXP uses acknowledgements and timeouts in order to ensure correct communication. CDXP is a subscription-based protocol, and uses notifications whenever data subscribed for has changed. When a timeout occurs on either the client or the server, a keep-alive message is sent that must be acknowledged.

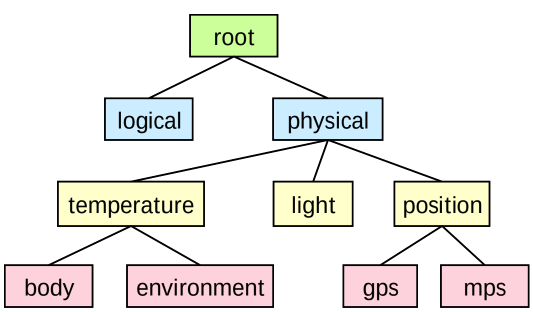

Fig. 4. Context Taxonomy

D. Context Servers

Context Servers collect and store context data from mobile devices via CDXP. Context Servers leverage the fact that they are stationary — not being limited in processing power, bandwidth, power consumption, etc. — to collocate a CDXP server with data-mining services and functionality for obtaining context information from various nodes in mobile networks, e.g.:

- BSC (resource parameters)
- HLR/VLR (GSM mobility)
- SGSN (GPRS mobility&session management)
- GGSN (AAA, IP address & routing).
- WLAN (SNR, Bandwidth, current AP, etc.)

Data-mining services in the Context Server synthesize context knowledge out of data received from these various sources and make predictions about when certain conditions are met. CDXP makes context knowledge available to Agents running in mobile devices.

E. Discussion

This system enables mobile stations to leverage context knowledge and use this knowledge to make more optimal choices between available wireless interfaces for communication between mobiles, as well as non-local decisions such as selection of a suitable neighbor as intermediary for sharing of content in a group of mobile users. For example, the Agent in the mobile can use context knowledge & predictions obtained from the Context Server to make decisions about services, such as deferring a service between peers, to a certain time & location, when the cost of transferring data is minimal and within the bounds of when this data should arrive (e.g., “context-aware Gnutella”).

VI. Conclusions

Combining the Agent and a SIP-UA (a Presence Agent that understands novel extensions for service discovery and has access to context knowledge via CDXP) is particularly important. We are in the position to create reconfigurable overlay service networks of context-aware entities. SIP allows us to build scalable solutions for large mobile networks and move service integration out to the end-points.

These early results focus on *infrastructure context* rather than *user context* (see further section I). Equally important, we examined how context can be used to prepare mobile users to for interaction with entities that they are likely to discover *prior* to their arrival in the area that these entities are located in. The importance of such an approach is that it allows us to *minimize* required user attention for interaction with communication infrastructure and instead *maximize* user attention for interacting with services that are relevant to the user's goals and situational context. We showed that interaction anticipation, the ability to ad-hoc include other discovered entities, we can enable spontaneously orchestrated and cooperative behavior of various entities (people, places, objects).

In our work we found (as reported in this paper) that geographical location (areas) and group membership are important scoping mechanisms, which will be necessary from a user interaction point of view (in order to be able to focus on relevance) but also from a communication point of view (in order to enforce effective communication patterns). Furthermore, this poses questions about managing policies. We conclude that the support outlined here for cooperative mobile behavior is vitally important for the creation of novel and affordable applications for mobile users in heterogeneous wireless infrastructure.

VII. Current Work

During the course of 2003 we will explore these issues and extend our prototype and testbed from the Pilot Project with sample applications that will provide us with evidence regarding these issues:

- Entity anticipator (PrePlug & Play) featuring cooperative anticipation of entities (incl. mapping of devices & services) in smart spaces, provide interactive guidance to users
- Explore emergent utility models (by inferring new uses from context knowledge), as well as remote invocation

In forthcoming publications we will report on empirical and theoretical results regarding the communication properties of the exchange of context data via these new protocols and to what extent the approach enables further optimization of network resources as well more optimal delivery of content according to the user's context. Furthermore, we will further examine modeling of context data and integration of our infrastructure with SIP Presence support [8].

VIII. Acknowledgments

This research is conducted at the Wireless Center of the Royal Institute of Technology, and funded by the Swedish Foundation of Strategic Research. I would like to acknowledge my co-researchers Fredrik Kilander, Martin Jonsson, & Li Wei, master thesis students Andreas Wennlund, Asim Jarrar, & Roberto Cascella, and finally professors Carl-Gustaf Jansson & Gerald Q. Maguire Jr. for their support and contributions to this research.

References

[1] Andreas Wennlund, "Context-aware Wearable Device for Reconfigurable Application Networks", Masters of Science thesis, Department of Microelectronics and Information Technology (IMIT), Royal Institute of Technology (KTH), Stockholm, Sweden, March 2003.

[2] Roberto Cascella, "Reconfigurable Application Networks through Peer Discovery and Handovers", forthcoming Masters of Science thesis, Department of Microelectronics and Information Technology (IMIT), Royal Institute of Technology (KTH), Stockholm, Sweden.

[3] Asim Jarrar, "Context Server support for Opportunistic and Adaptive Mobile communication", forthcoming Masters of Science thesis, Department of Microelectronics and Information Technology (IMIT), Royal Institute of Technology (KTH), Stockholm, Sweden.

[4] G. G. Richard III, "Service Advertisement and Discovery: Enabling Universal Device Cooperation," IEEE Internet Computing, vol. 4, no. 5, Sep./Oct. 2000, pp. 18–26.

[5] J. Rosenberg et al., "SIP: Session Initiation Protocol", IETF Networking Group, RFC 3261, June 2002

[6] A. B. Roach. Session Initiation Protocol SIP-Specific Event Notification. RFC 3265, IETF, June 2002.

[7] T. Kanter, "Adaptive Personal Mobile Communication, Service Architecture and Protocols", Doctoral Dissertation, Department of Microelectronics and Information Technology, Royal Institute of Technology (KTH), November 2001. Available at http://psi.verkstad.net/publications

[8] J. Rosenberg, "A Presence Event Package for the Session Initiation Protocol (SIP)", IETF SIMPLE Working Group, Internet Draft, January 31, 2003

[9] P. Bellavista, A. Corradi, C. Stefanelli, "The Ubiquitous Provisioning of Internet Services to Portable Devices", IEEE Pervasive Computing, IEEE Computer Society Press, Vol. 1, No. 3, pp. 81-87, July-Sep. 2002.

[10] A. Held, S. Buchholz and A. Schill, "Modeling of Context Information for Pervasive Computing Applications", Proc. of the 6th World Multiconference on Systemics, Cybernetics and Informatics (SCI2002), Orlando, FL, USA, July 14-18, 2002

[11] Mark T. Smith and Gerald Q. Maguire Jr., SmartBadge/BadgePad version 4, HP Labs and Royal Institute of Technology (KTH), 2002-08-12. http://www.it.kth.se/~maguire/badge4.html