

iClouds – Peer-to-Peer Information Sharing in Mobile Environments

Andreas Heinemann*, Jussi Kangasharju,
Fernando Lyardet, and Max Mühlhäuser

Telecooperation Group, Department of Computer Science
Darmstadt University of Technology
Alexanderstr. 6, 64283 Darmstadt, Germany
aheine@gkec.tu-darmstadt.de
{jussi, fernando, max}@tk.informatik.tu-darmstadt.de

Abstract. The future mobile and ubiquitous computing world will need new forms of information sharing and collaboration between people. In this paper we present *iClouds*, an architecture for spontaneous mobile user interaction, collaboration, and transparent data exchange. iClouds relies on wireless ad hoc peer-to-peer communications. We present the iClouds architecture and different communication models, which closely resemble familiar communication forms in the real world. We also design a hierarchical information structure for storing the information in iClouds. We present our prototype implementation of iClouds which runs on wireless-enabled PDAs.

1 Introduction

People living in the future ubiquitous computing world will need new ways to share information and interests as well as collaborate with each other. Given the success of wireless communications, such as mobile telephones, 802.11b, or Bluetooth, many of these activities will benefit from or rely on wireless ad hoc communications.

This paper presents the iClouds project in which we study the architectures and mechanisms required to support mobile user interaction, collaboration, and transparent data exchange. The iClouds project is part of the Mundo research activity, which we will present in the next section.

Our motivation behind iClouds can be expressed as follows: “Whenever there is a group of people, they may share a common goal or have a related motivation. Information of interest may be in possession of only a few of them.” The goal of iClouds is to make this information available to the whole group, based on individual user contribution, through peer-to-peer communications and data exchange.

* The author’s work was supported by the Deutsche Forschungsgemeinschaft (DFG) as part of the PhD program “Enabling Technologies for Electronic Commerce” at Darmstadt University of Technology

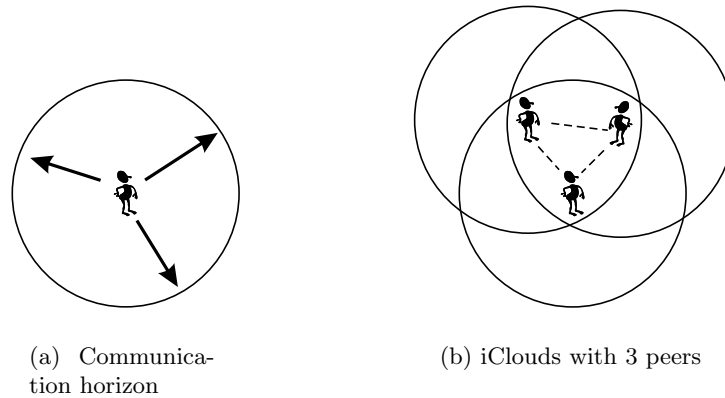


Fig. 1. Information clouds

Consider a person walking with a wireless-enabled PDA. The communication range of the wireless device defines a sphere around that node. We call this sphere or communication horizon an *information cloud* or *iCloud* (see Fig. 1(a)). In practice this will not be an ideal sphere due to radio signal interference with buildings and other structures. The limited communication range (a few hundred meters at most) is a desired property, because it allows for easy ad hoc meetings and collaboration. When several nodes come close together, as shown in Fig. 1(b), the devices can communicate with each other and exchange information depending on what information the users provide and need. This exchange happens automatically, without any need for direct user intervention.

We have identified several application scenarios in which iClouds is beneficial:

- **Local Information Acquisition.** Residents of a city publish information about their city which tourists are interested in. This could be information about sights, restaurants, or useful telephone numbers, such as taxi number, etc.
- **Common Goal Pursuit.** iClouds can bring people with common interests together to help them collaborate. For example, consider students in a classroom. Some students may have formed study groups and others might be interested in joining those groups. Students already in the groups could publish the groups and interested students could directly contact them and join the group.
As an alternative example, consider two (or more) customers planning to buy the same or a similar item. When they notice each other within communication range, they will have the opportunity to discuss the item, exchange experiences, and collaborate (e.g., ask for discount).
- **Advertisement and mCommerce.** A store can publish ads which are picked up by interested customers. These customers further pass the ads along to other interested users when they are away from the store, thus

increasing the reach of the ads. If any of the users who have received ads in this way actually make a purchase, the store could give a bonus to the person who passed the ad along. This bonus could be, for example, points or a discount on the next purchase.

This paper is organized as follows. Section 2 describes the underlying Mundo project. Section 3 presents the iClouds architecture and communication mechanisms. In Section 4 we describe our prototype implementation of iClouds. Section 5 discusses related work. Finally, Section 6 concludes the paper and presents directions for future work.

2 Mundo Overview

A major issue in a ubiquitous computing world is that of future personal devices in the dawning post-PC era. Our studies revealed the following key requirement for such a device: Nomadic users carry a (minimal) personally owned device, the ME (*Minimal Entity*). Our decision for one individually owned device rather than a publicly available infrastructure comes from the idea of *trust establishment*. As an increasing amount of computer based activities have (potential) legal impact and involve privacy issues, it is vital that users can trust the device carrying out such transactions. Trust establishment means not only physical ownership, but also functional properties. Such a device (i) holds and autonomously carries out functions on private keys and (ii) restricts the potential effect of network-originated functions and data.

Figure 2 shows an overview of our Mundo architecture. The word Mundo alludes to the Hispanic word “mundo”, meaning world; it can also be viewed as an acronym for *Mobile and Ubiquitous Networking via Distributed Overlays*. A more complete description can be found in [1]; below we will present an overview of the main components of Mundo, their interactions and relation to iClouds.

The ME is designed to be an ears-and-mouth device. A device meant to be always carried along must be as small as possible, typically even smaller than today’s cellular phones. Hands-and-eyes-free operation of the device is desirable both because of minimalization concerns (a ME device may shrink down to earplug-size a few years from now) and because of the possibility to have immediate access to the ME in most situations. Additionally, the demand for audio communication is ubiquitous, as demonstrated by the success of cell phones. In noise-sensible environments, such as meetings, head movement detection may replace a minimal audio command set. Sophisticated eyes-and-hands devices may always be used by means of association (see below).

We consider ME devices as the representation of their users in the digital world. The personal ME is the only entity always involved in the user’s activities. Given that the ME is a single point of failure, it is essential that the ME operates only on behalf of its owner in physical presence of that user. A stolen, malfunctioning, or out-dated ME must be easily replaceable by a new model or device.

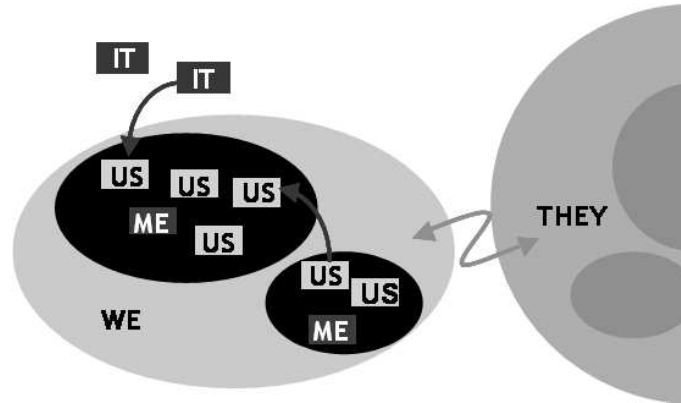


Fig. 2. Mundo architecture

Minimalization pressure will not permit feature-rich MEs in the near future. Hence, they must be able to connect to local devices such as memory, (co-)processors, displays, and network devices in an ad hoc manner. This goes beyond the traditional notion of service discovery in environments like Jini which basically exchange service APIs. We call this process association and we call such devices *ubiquitous associable objects* (US). Any US must permit its user to authorize information exchange between ME and US devices. Users must also have the possibility to mark an US as trustworthy. For privacy reasons, any personalization of an US must become unavailable if it is out of range of the user's ME device.

While USes are closely associated to the ME and therefore to the user, there will be numerous smart items that do not support association that would turn them into an US. Vending machines, goods equipped with radio frequency IDs, and landmarks with “what is” functionality are just a few examples. We define such *smart items* as ITs. An IT is any digital or real entity that has an identity and can communicate with the ME. Communication may be active or passive; memory and computation capabilities are optional.

In a ubiquitous setting, support of ad hoc networking is crucial. Mundo will feature so-called *wireless group environments* (WE). We expect ad hoc networking to be restricted to an area near to the user of a ME device, as connections with remote services will involve a non-ad hoc network infrastructure. The functionality of a WE is to bring together two or more personal environments consisting of a ME and arbitrary US entities each. Besides making connections between the devices possible, a ME must provide means for sharing and transferring hardware (e.g., US devices) and software/data artifacts between WE users.

The iClouds project presented in this paper is one example of a WE. The goal of iClouds is to enable sharing information contained in the ME or its associated US devices. Our focus in this paper is on pure information sharing; we do not consider sharing hardware between ME devices.

We regard overlay cells as the backbone infrastructure of Mundo. These *tele-cooperative hierarchical overlay cells* (THEY) connect users to the (non-local) world, deliver services, and data to the user. THEY support transparent data locations, for example, frequently used data may be cached on US devices. THEY offer cooperation between different physical networks, transparent to the users. Additionally, THEY must provide an internal model for cost, privacy, and accounting supporting cooperation with other networks. Because iClouds is based on ad hoc networking, it does not require the services of THEY.

3 iClouds System Description

iClouds devices are small mobile devices (like PDAs) with mobile communication support for a maximum of a few 100 meters; one example is a PDA with 802.11b support. There is no need for any central servers in the iClouds architecture; instead, each device is completely independent. In the Mundo context, an iCloud device would likely be an US that is associated with the user's ME. The US device would contain the information and the communications between users would happen either via their MEs or through the US devices (authorized by the respective MEs).

The diameter of the iClouds communication horizon (Fig. 1(a)) should not exceed a few hundred meters. We want to give iCloud users the option for spontaneous collaboration and when two iCloud users "see" each other, they should be within a short walking distance from each other (a couple of minutes at maximum). To allow for this easy collaboration, we specifically exclude multi-hop communication. Therefore, iClouds does not require any routing protocols; all communications happen directly between the concerned parties.

3.1 Data Structures, Communication Pattern and Information Exchange

The two most important data structures found on the iClouds device are two information lists (*iLists* for short):

- *iHave*-list (information have list or information goods)

The iHave list holds all the information the user wants to contribute to the iCloud. The items could be, for example, simple strings, or more complex entities expressed in XML. We discuss these issues in Section 3.2.
- *iWish*-list (information wish list or information needs)

In the iWish list, the user specifies what kind of information he is interested in. The exact format of the items on this list depend on the format of the entries on the iHave list. Typically they would be search patterns which are matched against the entries on the iHave lists of other users. Note that the items on the iWish list are more private, since they reflect the needs of the user, which the user may not want to disclose to others.

Table 1. Information Flow Semantics (from Alice’s point of view)

	pull (from Bob)	push (to Bob)
iHave-List	Standard search	Advertise
iWish-List	Active service inquiry	Active search

Each iClouds device periodically scans its vicinity to see if known nodes are still active and in communication range and also to see if any new nodes have appeared. Information about active nodes is stored in a data structure called *neighborhood*.

In the second stage, the iClouds devices align their information goods and needs. This is achieved by exchanging iLists. Items on the *iWish*-lists are matched against items on the *iHave*-lists. On a match, information items move from one iHave-list to the other.

For example, consider two iClouds users, Alice and Bob, who meet on the street. When their iClouds devices discover each other, they will exchange their iHave lists and match them locally against their iWish lists. If an item on Bob’s iHave list matches an item on Alice’s iWish list, her iClouds device will transfer that item onto her iHave list.

We have two main communication methods for transferring the iLists. Peers can either *pull* the iLists from other peers or they can *push* their own iLists to peers they encounter. In addition, either of these two operations is applicable to both lists, which gives us four distinct possibilities of communication. We summarize these possibilities, along with their real-world equivalents, in Table 1.

In each of the four cases shown in Table 1, the matching operation is always performed on the peer who receives the list (Alice’s peer in pull and Bob’s peer in push). Each of the four possible combinations corresponds to some interaction in the real world:

- *Standard search.* Alice pulls iHave-List from Bob.
This is the most natural communication pattern. Alice asks for the information stored on Bob’s device and performs a match against her information needs (specified in her iWish-List) on her device. We can also see the user as just passively “browsing” what is available.
- *Advertise.* Alice pushes her iHave-List to Bob.
This is a more direct approach. Alice gives her information goods straight to Bob and it’s up to Bob to match this against the things he is interested in. As an example, consider iClouds devices mounted on shopping mall doorways pushing advertisements onto customer devices when they enter the building.
- *Active service inquiry.* Alice pulls iWish-List from Bob.
This is best suited for shopping clerks. They learn at a very early stage, what their customers are interested. An example of this query could be: “*Can I help you, please show me what are you looking for?*”.




Fig. 3. Hierarchical data organization

In general, especially for privacy reasons and user acceptance, we believe it is a good design choice to leave the iWish-list on the iClouds device. Hence, this model of communication would likely be extremely rare in the real world.

- *Active search.* Alice pushes her iWish-List to Bob.

With active search, we model the natural “*I’m looking for X. Can you help me?*”. This is similar to the standard search mechanism, except that the user is actively searching for a particular item, whereas in the standard search the user is more passive.

3.2 Information Modeling and List Matching

A key issue in iClouds is a set of rules or constraints for the information items on the iHave-list, as well as for the search patterns stored on the iWish-list. For example, consider free text for both lists. A hypothetical item *Best Italian Restaurant on Market Street* will not match an information wish *Interested in Mediterranean Food*. Therefore the chances are close to zero, that any information is passed through iClouds, although one node holds information the other node is interested in.

To overcome this, we propose a hierarchical system, as shown in Fig. 3, similar to the organization of the Usenet, the product catalog of eBay, or the online catalog of Yahoo. Items are divided into categories (e.g., restaurants, hotels, etc.) which are further divided into subcategories (e.g., Italian, Indian, etc. restaurants). Note that entries may belong to several categories. For example, the Bombay Restaurant in Fig. 3 is listed under Indian and Vegetarian, because it offers both Indian and vegetarian food.

These hierarchies can be stored at a central server and can be downloaded to an iClouds device while it is in its cradle during power recharge or a normal sync operation. This system should be open for extension and categories could be moderated by different people to improve quality.

Hierarchies are well understood by users and, in addition, technologies, such as XML and XPATH, support the construction of hierarchically organized information and search very well. Fig. 4 shows an example of how a restaurant entry can be expressed in XML. The names of the most common tags, such as address or telephone, will need to be standardized. We will also need a standardized way for user extensions to the hierarchy. These are topics of our on-going work.

```

<restaurant>
  <name>Bombay Restaurant</name>
  <address>...</address>
  <telephone>...</telephone>
  <category>Vegetarian</category>
  <category>Indian</category>
</restaurant>

```

Fig. 4. XML Data Representation

3.3 User Notification

Collaboration regarding a common goal or a related motivation requires that users are notified by the iClouids device on a successful match. This can be a simple beep or a vibrating alarm as found in mobile phones today. In the Mundo architecture, the notification can also be sent to the user's ME. Recall that the ME device in Mundo is an ears-and-mouth device which supports basic audio communications. After notification, a voice based communication could be established for further information acquisition and possible arrangements for collaboration.

For each item (search pattern) on the iWish-list, the user can specify, whether a notification should be sent or not. This enables a user to differentiate between pure collect pattern, e.g. *Good Restaurants?*, and patterns that require some kind of action, e.g. talk to the other person.

4 Prototype

To gain more practical experiences with iClouids, we have built a first prototype and set up a testbed. The prototype runs on several Toshiba Pocket PC e740 (Windows CE) which are shipped with integrated 802.11b cards. Fig. 5 shows a screenshot. For the underlying link layer network connectivity, we run the PDAs in 802.11b ad hoc mode. The prototype was developed using the PersonalJava from the Java2 Micro Edition (J2ME).

The interface shown in Fig. 5 is relatively simple. It shows the iHave and iWish lists as well as other iClouids peers in the neighborhood. The user can mark the iLists as private; in this case other users are not able to retrieve them. There are also buttons for adding items to the lists and removing them.

Our information list data structures consist of strings. Currently, we have not yet implemented hierarchies. iList comparison is based on a simple substring matching function. A successful match will copy iHave-list items to new devices. The user is notified of this event by a beep from the device. This allows the user to check her updated information goods and plan further action.

We use a UDP based ping/pong mechanism for scanning the vicinity for new nodes. Periodically, a node sends a ping message to other nodes in its neighborhood and waits for a pong message. Upon receiving a pong, the new node is




Fig. 5. iClouds Pocket PC Prototype

added to the neighborhood. Otherwise, after a certain timeout, the node is removed from the neighborhood. When a node encounters a new node, it pulls the iHave list from the new node using an HTTP GET request. We have decided to use HTTP since it is a very simple and well-understood protocol and it is widely supported. This offers us the possibility of easily extending iClouds to standard web browsers or servers.

Using the prototype, we have learned that a PDA is not a good device for data input. Therefore, we believe that the iLists should be managed in a desktop application and the lists should be synced to the PDA during the normal sync operation.

5 Related work

The Proem Platform [2] targets very similar goals. The main difference to iClouds is that they focus on Personal Area Networks (PAN) for collaboration. We believe that it is fruitful to focus on a wider area (mobile networks that cover several 100 meters in diameter) and that it is not necessary to encounter communication partners physically for information exchange.

Spontaneous user collaboration was investigated by the Meme Tags project [3]. Meme Tags are small devices with LCD screen (two lines, about 30 characters), bidirectional infrared communication support, two buttons, and a knob for user control. As the information is only shown on the LCD screen and the transmission uses IR communications, Meme tags allow only for short-range, face-to-face communications. iClouds uses longer range communication methods, such as 802.11b, and allows for a much greater flexibility in information sharing and collaboration.

A key element to the success of iClouds is the desire of the users to contribute to the community. However, most systems that are based on user contribution,

have a certain amount of “freeriders”, that is, users who do not provide anything but only consume. Adar and Huberman [4] studied this freeriding phenomenon in Gnutella and found out that a large fraction of the users did not contribute to the community. We believe this is mainly because contributing to a file sharing network, like Gnutella, consumes a lot of end-user resources, especially bandwidth, and users who have less resources are unwilling to share them.

In contrast to this, contributing to iClouds would not require the users to sacrifice large amounts of their personal resources, and this is why we believe that the freeriding phenomenon would have less impact on iClouds. This is true for many systems which are based on small user contributions. The Usenet news is certainly one of the most prominent and successful systems. Another example that exploits small user contributions from a wider community is the *Vipul’s Razor* project [5]. This is a collaborative spam detection and filtering network. The idea is quite simple. Spam is not recognized by sophisticated scanning software, but by the individual user. This recognition is propagated to the Razor network, so that all users have the benefit that spam is detected earlier.

6 Conclusion

In this paper we have presented iClouds, an architecture for supporting spontaneous user interaction, collaboration, and transparent data exchange in a mobile ubiquitous computing context. iClouds achieves these goals through ad hoc peer-to-peer communications between wireless-enabled devices. iClouds supports many natural forms of interaction, such as browsing for information, searching, and advertising. We have developed a hierarchical information structure for storing and matching information in iClouds. We have also implemented a prototype of iClouds which runs on PDAs with wireless LAN cards.

As part of our future work, we will develop a model for specifying the iClouds information hierarchies and ways to extend them. In addition, we plan on investigating ontologies as a possible way to improve the matching between items on the iLists.

References

1. Hartl, A. *et al.*, Engineering Multimedia-Aware Personalized Ubiquitous Services. In IEEE MSE 2002, Newport Beach, CA, December 2002.
2. Kortuem, G. *et al.*, When Peer-to-Peer comes Face-to-Face: Collaborative Peer-to-Peer Computing in Mobile Ad-hoc Networks, International Conference on Peer-to-Peer Computing (P2P2001), Linköping, Sweden, August 2001.
3. Borovoy, R. *et al.*, Meme Tags and Community Mirrors: Moving from Conferences to Collaboration, in ACM CSCW, Seattle, November 1998.
4. Adar, E., Huberman, B. A.: Free Riding on Gnutella, Technical report, Internet Ecologies Area, Xerox Palo Alto Research Center, September 2000.
5. Vipul’s Razor web site, <http://http://razor.sourceforge.net>, (seen Feb. 2003)