

Supporting Persistent Social Groups in Ubiquitous Computing Environments Using Context-Aware Ephemeral Group Service

Bin Wang, John Bodily, and Sandeep K. S. Gupta
Department of Computer Science and Engineering
Ira A. Fulton School of Engineering at Arizona State University
Tempe, AZ 85287
{Bin.Wang, John.bodily, Sandeep.Gupta}@asu.edu

Abstract

In this paper, we analyze the role of the social group in a Ubiquitous Computing (UbiComp) environment as a source of contextual information. A model is presented to address the social group member's perceptions of how devices in a UbiComp environment should aid them in collaboration. Based on the model a distributed context-aware group membership management scheme is developed. We then present a prototype implementation for a context-aware ephemeral group membership management scheme, along with a sample application, and experimental results that demonstrate the feasibility of our system.

1. Introduction

A Ubiquitous Computing (UbiComp) environment contains devices that are embedded in the environment, and a set of transient wearable, and handheld devices that are being carried by users inside the environment. Devices are connected over wireless networks, and may possibly be connected to a fixed network infrastructure such as the Internet. Devices in such environments have serious resource constraints, and they form numerous webs of short-range low-power mobile ad-hoc networks to exchange information. A principal goal of a UbiComp environment is to provide computational support to help and mediate the users activities in a natural, and non-intrusive way [17, 7]. This requires designs to be user-centric, and not simply device-centric. Most research has focused on supporting users on an individual basis in UbiComp environments. It is common for users to be members of some social group. Thus, in order for interactions between the user and the computing environment to be seamless, support must exist for the collaborative efforts of groups as well.

1.1. Social Group

The definition of a social group that is used for this paper is defined in [11] as follows: "A number of individuals, defined by formal or informal criteria of membership, who share a feeling of unity or are bound together in relatively stable patterns of interaction". This definition stresses two properties of the social group; the first is that people in a social group must feel some sense of membership. Second, social groups have a set of predictable collaborative interactions that group members engage in with one another.

In order for a person to be a member of a group there must be consensus between group members to determine their inclusion. Membership is not decided solely at the individual level. The group, as a whole must recognize it. For example, in a software engineering class students are often assigned to groups. In these groups there is usually a provision made for the group to decide whether or not a student will persist as a member of the group. Thus, a student must identify themselves as a member of the team, but also the team must continually recognize the student as a member.

An interaction indicates some collaborative activity that requires one or more users to perform a set of actions towards the completion of a common goal. Because a group is formed for a specific purpose, the collaborative interactions become predictable. For example, a team of software engineering students will have a predictable set of collaborative tasks as they work on gathering requirements, designing, testing and evaluating software. The students will also have regular meetings to discuss progress, and plan for future milestones. It is from the concept of membership, and stable interactions in a social group that we derive our concept of "persistence" in the social group.

In general when a person is identified as a member of a group they're considered a member of that group for a considerable amount of time. We refer to this as the "lifetime" of that person's group membership. Some member-

ships such as being a family member have a lifetime that starts when one is born and ends one dies. The membership lifetime in other groups may last only a few months, such as a school project group. However, when compared to the time it takes for a user to complete an individual session on a computer, the lifetime of a social group membership is generally much longer than the time taken to complete the computing session. So membership in a group is a source of persistent state data that cuts across computing session boundaries. From this point of view, social groups act as a source of persistent contextual information that should be leveraged by computing applications in order to seamlessly support users in a Ubicomp environment.

1.2. Social Context

Both location and identity contribute to the characterization of a situation that a user experiences but they are not the only sources of context present in an environment. A broader spectrum of contexts has been discussed in regards to context-aware computing in [16]. This paper is more concerned with social contexts:

Definition 1 Social Context: *Information relevant to the characterization of a situation that influences the interactions of one user with one or more other users.*

Social mores and norms are an example of social contexts. These influence the proper use of technology in a given environment. For example, when at a movie theatre it is impolite to allow a cellular telephone to ring during a showing. No law states that your phone cannot ring during a movie, it is just a social norm to consider it rude.

A large part of multi-user software system design revolves around determining the functional scope of the system by capturing the social contexts that govern how users are allowed to use the system as a tool. When building such a system the intended usage of the system is often examined using use-cases. Then based on certain social roles that are designated to a user by a governing authority (generally a business) functionality is restricted to a subset of the total system functionality and controlled using access rules.

It was established in [15] that in order for software to be unobtrusive to the user in a Ubicomp environment that it must become context-aware. As illustrated by the previous section, social contexts have played a large part in traditional system design. However, filtering the content and functionality available to a user through a set of persistent static access rules and roles does not make a system context-aware. In order for a system to be context-aware it must interface with both the user and the environment to dynamically sense the context and adapt its functionality in order to best serve the user for a given situation. Social contexts are intangible in that you cannot build a sensor to register

readings for a social more. Is there a way to leverage other contexts to characterize social contexts and detect them? An example follows that will be used to illustrate the challenges associated with overcoming this goal.

Example 1 *Consider a group of students taking a software engineering course that are involved in regular meetings to peer review reports. Assume that each user carries a PDA that runs software allowing the users to participate in some computer mediated collaborative review process. In order for this to be ubiquitous the software should support this in environments with either fixed network infrastructure or in ad-hoc environments. Now, for a document peer review we must assume that the software supports the roles of reviewer, recorder, author, and moderator. When all members are in attendance at the designated location, at the correct time the application should start the peer review process. At some point during the review if the author and a reviewer choose to leave to review the raw data for the results presented in the report they should be able to continue working with one-another independent of the other group members. Similarly when they return to the peer review they should be integrated into the group as a whole again seamlessly. Then once the review process is over the software should end the review process and send the meeting minutes from the review to everyone.*

1.3. Challenges of Supporting Groups in Ubicomp Environments

The main challenge illustrated by Example 1 is the seamless support of the group's interactions once membership has been established. The social context we are interested in detecting is the fact that a peer review is in progress. As we cannot sense the actual event of "Peer Review" the only option is to infer the occurrence of the event somehow using other contexts. The following challenges are specific to finding a solution for this problem:

1. **Determining when an interaction is going to occur -** During the lifetime of a user's group membership several collaborative group interactions will occur which will require communication links to be established between devices owned by the group. These interactions occur based on some social context. A social context must first be represented in a format that can be used to perform computation in order to detect the occurrence of the collaborative interaction that this context denotes. Furthermore, computation should be supported in a repeatable manner with minimal demands on the end-user.
2. **Determining when an interaction is going to end -** Most interactions will have an end, which is also denoted by some social context that needs to be detected

and used to do processing specific to supporting the end of a group interaction. Similar to detecting the beginning of an interaction, we also need a way to represent the social contexts that characterize the end of a group interaction. Again, this should be supported with minimal demands on the end-user.

3. **Supporting interactions in environments with varying network capabilities** - For a truly ubiquitous solution, support offered to the group should not be dependent on the networking capabilities of the environment. A solution should be able to function in both fixed and ad-hoc network environments. In order to provide a solution that is well suited to varying network capabilities a distributed solution is preferred. The solution must also allow a user or sub-group to leave a group due to becoming disconnected from the other members, and then rejoin once a connection can be re-established.

The rest of the paper is organized as follows. In Section 2, we describe our system model and a group model for Ubicomp environments, which focuses on demonstrating and presenting the relationship between two key groups: user group and device group. Section 3 describes a distributed Context-Aware Ephemeral Group (CAEG) Membership Management service. Section 4 presents a sample Group Chat application to demonstrate how the group model can be combined with CAEG to support discussion between users over ad hoc networks. In Section 5, we present the testbed which we have used to implement the CAEG service and the Group Chat application. Then some related work is discussed and used as the basis for comparison in Section 6. Finally we give some concluding remarks in Section 7.

2. System and Group Model

2.1. System Model

Our system model is designed around four primary assumptions. First we assume that each user will carry one or more Ubicomp devices that can communicate over a wireless network. Second, we assume that the users collaborate by using distributed peer-to-peer application software over a mobile ad hoc network with provisions for multi-hop message routing. Three, we assume that there is no Byzantine failures in the system. This assumption considerably simplifies our design. However, this restricts the application of our system to those scenarios where devices (users) trust each other. Four, we assume that the system clocks at individual devices are synchronized with each other within the tolerance desired by the application. Some applications may have very stringent synchronization requirement while oth-

ers may not require any synchronization at all. The problem of how to achieve synchronization in a distributed system is a well studied problem and we do not address these issues here in this paper.

2.2. Group Model

The system must manage the user's view of the social group in terms of membership criteria and collaborative interactions in order to support the user's perception of his/her social groups in a manner that minimizes human distraction [6]. We assume that consensus must somehow be achieved to decide on assessing the candidacy of a member for a group, and validating their membership in the group. As discussed in Section 1 this is a non-trivial task, and it is necessary to include the user in decisions about group membership. Similarly, discovering a social group to join is a non-trivial task, and may require user intervention. Some groups may openly publish their existence and allow new users to join freely. Other groups may wish to keep their existence a secret and protect the names of their members. The social factors involved in selection of membership for groups is very complex. Feedback from the user must be obtained, and full system automation is not desirable.

In order for the user to perceive that the system shares their notions about a social group, the system should provide unobtrusive support for the general routines (such as devices setup and group member discovery) that the group carries out. Furthermore, support for these routines should last for the entire lifetime of the user's membership with the group, and in order to reduce distraction should minimize effort needed by the end-user. As in traditional system design, limitations of the systems functional scope can be controlled by tailoring functionality to the roles that are present in a social group. We now present a model that adheres to the above mentioned criteria.

Based on the challenges discussed in Section 1 we have developed a conceptual model that primarily addresses the need to support group membership and routines with minimum system set-up. The fundamental goal of our approach is to support the social group with minimal cognitive effort expended on actual system operation, and to make users perceive that a device is a group collaboration tool. First the definition of a device context is presented. Next, the conceptual group model will be described. Finally a description of how the group model addresses Challenges 1 and 2 from Section 1 is provided.

Definition 2 Device Context: *Any detectable and relevant attribute of a device, its interaction with external devices, and/or its surrounding environment at an instant of time.*

For example, device context includes location of device, available memory of device, etc.

A description of our approach to developing a group model is now given. The basic premise of our model revolves around two different conceptual groups: User Group and Device Group.

User Group: The User Group is a social group that the end-user interacts with as discussed in Section 1.

Our model is based on the fact that user groups are expected to have stable patterns of interactions. This paper refers to these regular and repetitive collaborative interactions as a user group session:

Definition 3 *The time period for which the members of a user group are actively participating in interactions with one another is known as the User Group Session.*

In Example 1, a user-group session consists of a report peer review. The report review happens at a regular and reoccurring time, and for approximately the same duration each time.

The second group in our model is known as a device group. According to our system model we assume that an end-user will be carrying one or several computing devices that will be used to collaboratively interact with the devices of other group members. Logically a user may view these devices as a group that forms whenever the group needs to collaborate. This is referred to as a device group and is defined as follows:

Device Group: A device group is one or more devices that form a complete unit in composition with the purpose of collectively supporting some computational tasks on behalf of the user-group. The duration for which the device-group forms is determined by the user group session. However, not every device is a candidate for each device group. This implies that there must be a method for selecting the devices which become members of the device-group. The Device Context is used in this case to facilitate member selection for a device group, and in Example 1 we see that the Device Context can be used to allow only the devices owned by the team members to take part in the group. By giving the device group functionality to select the correct device members based on social context the end-user is given the perception that the device-group supports their understanding of social situations.

Figure 1 shows the overall relationship between the User Group, and Device Group. Notice that in this case the membership of a user in a user group is based on social contexts, whereas membership in a device group is characterized by the device context. The dotted line indicates that the user will develop the perception that the device group as a whole is interpreting the social context in order to support the user group. In order to provide this perception to the user each device must support device group formation based on the device contexts used to characterize the social contexts that decide when a group collaborative effort

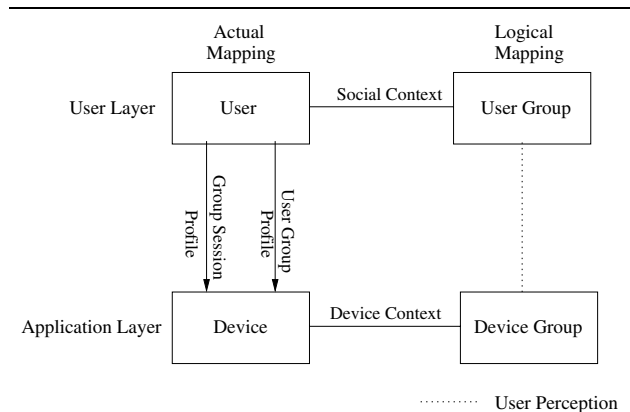


Figure 1. Multi-layer group model in Ubicomp environment.

begins and ends. As shown in Figure 1 our model uses profiles to accomplish this. A profile is nothing more than a file that resides on a device that is used to specify values of context. The values in the profile are used to trigger the formation of a device group, or provide informative social context information to an application in order to mediate and support both group interactions, and functionality specific to the role(s) of the end-user. The model uses the following two profiles as shown in Figure 1:

User Group Profile (UGP): This profile provides a representation of social contexts that are relevant to a user group, and are used by applications at a device that support collaboration to facilitate the specific actions required for a user's role or preferences. The User Group Profile includes a unique user identifier, a group identifier, the user's preferences, a group purpose and the role of the user in the group. In Example 1, each user will have a specialized function in the group such as reviewer, recorder and author.

Group Session Profile (GSP): The Group Session Profile is used to characterize some specifics that can be used to represent the contexts which define the beginning and ending of a collaborative interaction between group members. The profile should at least include the time period of the group session, the location where the group session will occur. In some cases be an "anywhere" location, and the interval of the group session. In Example 1, the group session profile is used to define the time period and location of a meeting for reviewing reports.

Now that all of the components of the model have been introduced we can analyze the model from the perspective of how it uses context in two ways. First to provide support for user group routines, and second to support the user's social roles and preferences.

Based on the group model we can see that the user has a perception of identity which is tied to the group. In particular this identity has a lifetime associated with it. In Exam-




Figure 2. Duration of groups in Ubicomp environment.

ple 1, the lifetime of the user's membership with the group lasts from the time it is assigned by the instructor to the time that the semester ends. During this time the group will have several routine user group sessions that need to be supported by device group formation. Context information represented in the Group Session Profile (GSP) is used by the device to form the device group that supports a user group session.

Figure 2 shows a graphical representation of times that the groups from the group model are being formed, and dispersed as viewed by a group member over time. If a user is a member of the User Group A during the time (T_A, T'_A) shown in Figure 2, they will participate in m routinely held user group sessions during the times $(t_{A_1}, t'_{A_1}), (t_{A_2}, t'_{A_2}), \dots, (t_{A_m}, t'_{A_m})$. Each device group A_i is shown in the figure as being active from the time it is formed, to the time it is dispersed, during an interval (t_{A_i}, t'_{A_i}) . Thus the context representation from the Group Session Profile after initial creation will be used m times to form a device group for each of the m user group sessions. The Context-Aware Ephemeral Group Membership Management described in Section 3 is used to form the device group.

We now examine how the model uses social contexts represented in the User Group Profile (UGP) to support the user group. As previously stated, the UGP is an input used by the devices to support collaboration between members in the user group. More specifically in Example 1 we assumed the existence of some distributed peer-to-peer application software to support document peer reviews, the UGP would be used as input for that application. Based on the user's preferences and through some pre-defined roles that are supported by the application the software will directly support the user's expectations of what interactions they should be allowed to perform with the group. This provides the user with a perception that the applications understand their needs and purpose in the group. In Example 1, a possible role that could be assigned to a user in the team should be moderator, and hence in the UGP for that person they will be described as having the social context of moderator. The document peer review application should use the information in the UGP to provide the appropriate functionality to the user.

A device group is a collaborative tool used to support a social group. Social contexts determine the inclusion of a user in a group, and the types of collaborative activities that a group performs. UGP is an abstraction of the social context that are relevant to a user group. The GSP specifies device contexts that describe when a social context indicates a collaborative interaction between group members will begin and end. The UGP contains information about social contexts relative to an individual member of the group such as their role in the group. By using the GSP to construct a group membership management service, and the UGP to customize application functionality the end-user is given the perspective that the device group interprets social contexts, and supports their user group. In this paper, we assume that all of the collaborative interactions in a group can be abstracted into a GSP and that contexts included in the GSP can be detected by devices. For example, after an exciting game of football, if the fans of the winning team decide to celebrate together after the game, their group celebration is a user group session. However, the beginning and ending of this user group session can not be abstracted into a GSP. Handling such spontaneous user group sessions is out of the scope of this paper.

3. Context-Aware Ephemeral Group Membership Management

In this section we study how to form device groups over ad hoc networks by using context to characterize the event associated with GSP in a distributed manner. Because device groups in a Ubicomp environment are ephemeral, they should be facilitated by a non-deterministic process group[10]. A group is defined to be deterministic if each member should receive and act on a request. Deterministic process groups in general require strong co-ordination and synchronization among the group members. As opposed to this, a Non-deterministic process group is used to implement distributed applications that do not require strong consistency between peers, and hence the overhead expended to manage the group will be more minimalist. Based on Challenge 3, we design CAEG membership management in a distributed manner, which is similar to [3, 9].

As we discussed in Section 2.2, the User Group Profile describes the function of the group and the role of the user in the user group. The UGP is a tuple $UGP = \{GID, GroupPurpose, UserRole\}$ where the fields are defined as follows:

- **GID:** *the Group Identifier is used to uniquely identify a Group on each device.*
- **GroupPurpose:** *the GroupPurpose describes the function of the group, for example "report peer review" or "discussion"*

- **UserRole:** a User Role describes the role of the user in the group, such as “recorder” or “author”.

The Group Session Profile (GSP) also introduced in Section 2.2 describes the session of a user group. For the sake of brevity, we assume that there are only two kinds of context (location and time) in GSP, although it would be easy to incorporate other kinds of context in GSP. $GSP = \{GID, Location, StartTime, FinishTime\}$.

- *StartTime, FinishTime:* When the CAEG starts and finishes.
- *Location:* The geographical area in which group members are located.

Using the GSP, a device group will be formed and dissolved. Actions at each device in a device group are based on the content of the Group Session Profiles at that device. For example, a user belongs to User Group A in Figure 2. User Group A has m device group sessions. And the user uses device d at device group session A_1 . Then the lifetime of d 's membership in A_1 is based on the duration for which d 's device context matches with the GSP for A_1 . At time t_1 ($t_{A_1} \leq t_1 < t'_{A_1}$), if d 's device context matches the GSP, d is able to become a member of device group A_1 . Later at time t_2 ($t_{A_1} < t_2 \leq t'_{A_1}$), d 's device context may no longer match the GSP resulting in d 's deletion from A_1 . And at time t'_{A_1} , the device group A_1 will disperse.

Assuming every device has already obtained a valid User Group Profile, a device group is then needed to support the user group session. The function of CAEG Membership Management Service is to automate the formation and disbanding of device groups for user group sessions by using the following steps:

- 1) **GSP Initiation and Distribution:** The group initiator defines the GSP and distributes it to the anticipated group members.
- 2) **GSP Registration:** Once the group session profile is transmitted and received, an application running on the device is notified about the arrival of the new GSP. The application will then be used to decide whether or not to accept the GSP. If the GSP is rejected, it will be discarded. Otherwise, if the GSP is accepted the application will register the GSP with the CAEG Membership Management Service. A GSP also contains a membership list that consists of the group member's identifiers. Initially, the GSP's membership list is Null. A GSP with a non-null member list is called an *active* GSP, and several active GSPs can exist on a single device to support various device groups simultaneously.
- 3) **Device Group Activation and Member Addition/Deletion:** A device gets data for the contexts that are relevant to the fields of a GSP from sensors. The data from the sensors is then used as input

to the CAEG Membership Management Service. The context values that are captured by the CAEG Membership Management Service are known as the *CAEG Context*, where the $CAEG\ Context = \{time, location\}$. The CAEG Membership Management Service monitors the *CAEG Context* and takes the following context-dependent actions:

If the *CAEG Context* C matches the contexts specified in the Group Session Profile gsp , i.e. $gsp.StartTime \leq C.time \leq gsp.FinishTime$ and $C.location$ is currently $gsp.location$, it adds the device identifier to the GSP's membership list.

If the CAEG Context no longer matches the Group Session Profile, the device identifier will be deleted from the GSP's membership list.

- 4) **Group View Maintenance:** A group view maintenance protocol for devices to exchange a *Group View* (GV) with each other in mobile ad hoc network has been developed. To deal with fault tolerance issues in group view maintenance, our protocol guarantees view accuracy and view completeness (The group view correctly expresses the group membership in all kinds of scenarios of network connection). Because the space limitation, we do not give the details of group view maintenance protocol in this paper. A device's Group View includes the GID from the GSP and the GSP's current group membership list¹. Using this protocol each device will know the Group Membership List of all devices in its vicinity, and according to those Group Views a device will update its own group view.




Figure 3. Architectural overview of CAEG membership management.

Figure 3 shows the architecture of CAEG Membership Management. There are three modules in CAEG Membership Management Service: Context Processor, Group Profile (UGP and GSP) Processor, and Group View Maintenance.

¹ A list of currently active and connected group members from the device point of view is the group view of the device group [4]. So a single device's group view need not be a global group view.

- a) The Group Profile Processor stores the UGPs and GSPs. Using Group Profile Processor, application can change to the UGPs and GSPs.
- b) The Context Processor abstracts the context that it obtains from sensors;
- c) The Context Processor passes the *CAEG context* to the Group Profile Processor; By comparing the *CAEG context* with the GSP, the Group Profile Processor determines whether or not to activate/deactivate the group session profile.
- d) After a GSP is activated, the Group View Maintenance protocol will exchange the group view with other CAEG services running on remote peer devices.
- e) If the current group view changes, the CAEG Membership Management Service will notify applications about the new group view.

To provide flexibility to users, CAEG Membership Management Service allows users to dynamically make changes to GSP. In Example 1, if the meeting of reviewing report is postponed by an hour, users and/or applications can increase *StartTime* and *FinishTime* of the GSP at their individual devices by one hour. This will ensure that the device group formation will be also be postponed for an hour.

4. Group Chat

This section will demonstrate how the User Group concept can be used to implement an actual application. The application in this case is a chat client named GroupChat for use over ad hoc networks. Group collaboration is supported by letting users communicate, and exchange files with one-another in settings where infrastructure support is unavailable. Lack of infrastructure support necessitates the need to design the application in a peer-to-peer manner with no reliance on a central server. The chat application we discuss uses the CAEG service, and the UGP to establish groups for the purpose of communication, and file sharing.

Users are organized for GroupChat into social groups of users, which wish to communicate with each other. The groups used in GroupChat are analogous to a buddy list, or a chat channel. A user is a member of a social group if they have a valid group session profile, and user group profile on their device for that group. Two basic roles exist in the GroupChat application; a role of “chat moderator” who initiates the group, and the role of “chat user”. The basic interactions the application associates to users within the group with the role “chat user” are one to many communication, one to one communication, and file transfer. Additionally, an interaction must exist to view a user profile for each user in the group. Users with the role “chat moderator” should have access to all of the interactions provided to the “chat user” role plus the ability to perform “mute”, and

“unmute” interactions. The GroupChat application contains four major components as described below, and depicted in Figure 4:

ChatGroup Manager: The ChatGroup manager updates the application data needed to characterize the state of activated groups. It uses the User Group Profile to determine which interactions the end-user should be allowed to perform on a per group basis, and Group Session to determine the user group sessions. All text sent to, or received from a group is buffered here for retrieval by other GroupChat components.

GroupUpdate Processor: This component is responsible for registering the application with the CAEG service. It also processes notifications sent from the CAEG about changes to the the current set of active user groups, and to changes in the membership views for those groups. The results of processed notifications are used to update the ChatGroup Manager’s state.

ChatMessage Processor: The ChatMessage processor interacts with the ChatGroup manager and the CAEG service to send, and receive messages, or files destined for a remote-host in a group. Messages still needing to be sent are extracted from the ChatGroup Manager and passed to the CAEG service, and vice-versa.

Application GUI: The application GUI provides the user with the functionality to carry out the role-based interactions that they are allowed to execute. It must also display to the end-user the current set of activated groups, the membership views for those groups, as well as text messages, and the status of file transfers.

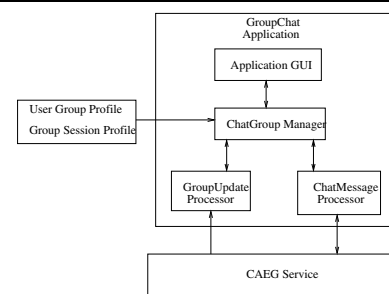


Figure 4. Architecture of the GroupChat application.

The GroupChat application has been developed as a sample application in our research testbed that is discussed in the next section. Figure 5 shows a screen-shot of the prototype GroupChat application.

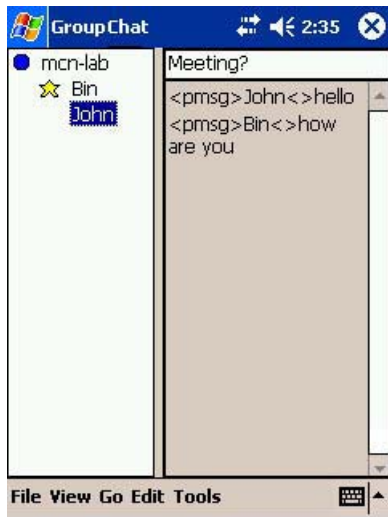


Figure 5. Screen-shot of the prototype GroupChat application.

5. Prototype of CAEG Service and GroupChat Application Implementation

To evaluate the CAEG service, we are currently implementing a Smart Classroom [18]. The test bed will facilitate different activities leading to efficient teaching and collaboration in a classroom. Smart Classroom monitors the classroom context (location of students or instructor, noise, light, mobility) and uses the context to trigger communication activity among students and instructor to facilitate collaborative learning [18]. Also to deal with heterogeneous environment issues, a Reconfigurable Context-Sensitive Middleware (RCSM) has been developed [19].

5.1. Test Bed Configuration

In our current design, each node in the test bed has the configuration shown in [19]. The configuration consists of the following components: Casio E-200 PDAs, Trez Electronic USB-compatible Xilinx Spartan II FPGA boards, noise, light, and motion sensors and location-tracking components. Each PDA uses an Intel Strong Arm 1110 with 206 MHz clock speed CPU. Each PDA had Flash ROM and RAM of 32 MB and 64 MB respectively. Each PDA was also equipped with a D-LINK Air DCF-660W Compact Flash 802.11b adapter. These adapters were configured in ad hoc network configurations. As such, no infrastructure, such as an access point, was necessary to provide the communication support.

We have used Microsoft Embedded Visual Studio to develop the CAEG service, and GroupChat application for the

Windows CE operating system. A screen-shot of the chat application as captured from the PDA is shown in Figure 5. The size of the executable programs for the prototype of the CAEG service, and GroupChat application is 35KB and 84KB, respectively.

5.2. Performance Results of CAEG Membership Management Service

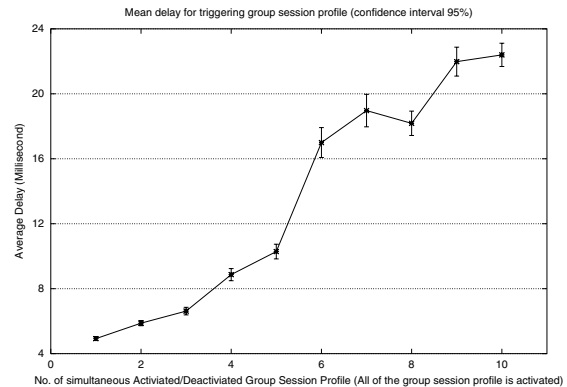


Figure 6. Mean delay for triggering GSP (All groups are triggered).

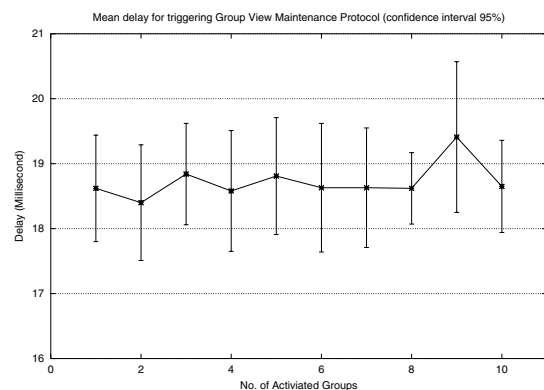


Figure 7. Mean delay for triggering group view maintenance protocol.

We use two performance metrics to denote the latency of CAEG Membership Management: The delay of activating a GSP and the delay of triggering the group view maintenance protocol. The delay of triggering the GSP is the time period from when the CAEG Membership Management obtains raw context data to the time when the corresponding

GSPs are activated/deactivated. The delay of triggering the group view maintenance protocol is the time period from when a GSP is activated to the time when the first group view is sent out.

Figure 6 illustrates the delay of triggering the GSP. In Figure 6, we notice the correlation between the number of GSPs being maintained by the CAEG Membership Management, and an increase in delay. Figure 7 shows the delay of triggering the group view maintenance protocol. Notice that the delay remains stable even if the number of activated group profiles increases. Based on an observed overall end-to-end delay of approximately 1 second in our testbed, we can see that the additional overhead introduced by CAEG Membership Management produces an imperceptible amount of delay.

5.3. Lessons from Experiments

From the experimentation that was conducted with the test-bed software and the CAEG service, we realized the importance of factoring resource optimizations that encompass the whole, rather than just the parts, into design. It is especially important when applications must run on resource-constrained devices such as PDAs used in our experiment. With no infrastructure support to share the burden of facilitating communication or computation the importance of factoring energy-efficiency and fairness into designs becomes very apparent. For example when implementing the CAEG group view maintenance protocol the design could have been carried out for a leader based protocol where the leader node once elected beacons requests to neighboring nodes to send in a message to indicate their presence, and then once a certain time interval elapses the leader node would send the view out to the other nodes that it received replies from. While this solution is valid, it does not consider how quickly the leader will lose power, and hence when it must be shutdown the group view will be lost.

6. Related Work

Few studies have been done for group communication in a Ubicomp environment. Prakash and Baldoni give a simple architecture for group communication in mobile systems [12]. In their architecture, a proximity layer below a group membership layer is introduced. They use a three-round group construction protocol to construct a group. The core idea is to create a group of all the nodes that are within a given distance D from the group initiator. They use a flooding algorithm to discover the nodes located in the area. However, there is no support for disconnection and context processing in the architecture, and a node can only be a member of a single group. Killijian et al. in [9] describe a location-aware group, which is similar to [12]. The group

membership of mobile devices in this scheme depends on the location of the devices, but they do not make provisions for ad hoc networking, or any contexts except location. Roman et al. proposed a mobile group, which is defined by a safe distance [14]. Only the nodes that are within the safe distance of its nearest neighbor can be considered as a group member. This notion of a group is very restrictive and does not deal with the general understanding of group communication; however, it is useful to implement disconnection awareness and network partition anticipation.

Briesemeister et al. [3] presents a localized group membership service that maintains the group membership status of neighbor nodes (Neighbor nodes are the nodes within a node's maximum transmission range) and proves the correctness of their implementations. Both Babaoglu et al. [2] and Fekete et al. [5] study partition-able group communication services to support continuing operations of "partition-aware" applications without blocking in multiple concurrent partitions.

Context-aware applications adapt their behavior in order to better meet the needs of the end-user in their current environment, while trying to make these adaptations as unobtrusive as possible. Applications such as the CyberGuide [1] from Georgia Tech use the context of location to provide relevant and useful information to tourist about nearby attractions, or displays inside a museum. Portable Help Desk (PHD) [6] developed at Carnegie Mellon University uses both spatial and temporal context to provide several services to students when they are on campus. This includes map services that display nearby available resources which can be computing and non-computing related, such as nearby printers, and events on campus. PHD can also be used to locate other individuals on campus whom may be a friend or a class team member. Context-awareness has been blended with instant messenger and chat applications in the past by the Hubbub [8] project at AT&T, and by the ConChat [13] application developed at University of Illinois at Urbana-Champaign. The Hubbub project uses the context of sound, to foster opportunistic interactions with users on the buddy list by playing a unique sound for each user when they become available. ConChat uses context-awareness to aid parties in determining the overall context in which an on-line chat is occurring, by displaying the context present in the room where a remote user is chatting from. Neither Hubbub, or ConChat support messaging without support from a centralized server.

7. Conclusions

In this paper we have presented a model for supporting social groups in a Ubicomp environment. First, we explored the properties of the social group, and defined what is meant by a "Social Context". From this we determined that the key

areas needing support for Social Groups in a Ubicomp environment are: support for the persistent membership of a user in a social group, and support for automating the collaborative routines that take place between group members.

Next, a conceptual model was constructed showing the relationship between a social group, and the computing devices used to support those groups. The model also illustrates how to use a device that is context-aware to form a device group based on contexts such as location, and time in order to support group collaboration, and provide the user with the perception that the device group is interpreting social contexts to support them. Third, we described a distributed CAEG membership management scheme that facilitates the formation of the context-aware ephemeral device groups discussed by our conceptual model.

Then a description of a sample application that demonstrated the overall concepts of our approach was given. We also discussed the details of our testbed, the primary focus of our usability testing, and the experimental results that were produced by that testbed. We also discussed the lessons learned from conducting our experimentation and discussed some related work.

In future, we plan to conduct usability tests to demonstrate how much user distraction can be reduced by using CAEG service. The focus of our usability testing will be to determine how well the implementation that we have developed addresses the challenges in the Section 1. Users will be asked to complete a set of test scripts involving a scenario and some questions targeting their experience using the system, as well as listing any improvements that may be made. The preferred user for our testing will have a substantial understanding of the mechanics of using a PDA, and can have either a technical or non-technical background. In particular the scenarios are being designed to quantify the system learning curve, and the amount of time saved by all members in the group due to having membership automated using the CAEG Service.

8. Acknowledgments

This work is supported by National Science Foundation (NSF) grant ANI-0123980. We thank the anonymous reviewers and Crystal West for their helpful comments for improving the quality of this paper.

References

- [1] G. D. Abowd, C. G. Atkinson, J. Hong, S. Long, R. Kooper, and M. Pinkerton. Cyberguide: a mobile context-aware tour guide. *Wireless Networks*, 3(5):421–433, 1997.
- [2] O. Babaoglu, R. Davoli, and A. Montresor. Group communication in partitionable systems: Specification and algorithms. *IEEE Transactions on Software Engineering*, 27(4):308–336, Apr. 2001.
- [3] L. Briesemeister and G. Hommel. Localized group membership service for ad hoc networks. In *International Workshop on Ad Hoc Networking (IWAHN)*, pages 94–100, Vancouver, British Columbia, Canada, Aug. 2002.
- [4] G. Chockler, I. Keidar, and R. Vitenberg. Group communication specifications: a comprehensive study. *ACM Computing Surveys*, 33(4):427–469, 2001.
- [5] A. Fekete, N. Lynch, and A. Shvartsma. Specifying and using a partitionable group communication service. *ACM Transactions on Computer Systems (TOCS)*, 19(2):171–216, 2001.
- [6] D. Garlan, D. P. Siewiorek, A. Smailagic, and P. Steenkiste. Project aura: Toward distraction-free pervasive computing. *IEEE Pervasive Computing*, 1(2):22–31, 2002.
- [7] S. K. S. Gupta, W. C. Lee, A. Purakayastha, and P. K. Sriyani. An overview of pervasive computing. *IEEE Personal Communications*, 8(4):8–9, Aug. 2001.
- [8] E. Isaacs, A. Walendowski, and D. Ranganathan. Mobile instant messaging through hubbub. *Communications of the ACM*, 45(9):68–72, 2002.
- [9] M.-O. Killijian, R. Cunningham, R. Meier, L. Mazare, and V. Cahill. Towards group communication for mobile participants. In *1st ACM Workshop on Principles of Mobile Computing (POMC)*, Newport, Rhode Island, Aug. 2001.
- [10] L. Liang, S. Chanson, and G. W. Neufeld. Process groups and group communication: classification and requirements. *IEEE Computer*, 23(2):56–68, Feb. 1990.
- [11] G. Marshall and et al. *The Concise Oxford Dictionary of Sociology*. Oxford University Press, Great Britain, 1994.
- [12] R. Prakash and R. Baldoni. Architecture for group communication in mobile systems. In *17th IEEE Symposium on Reliable Distributed Systems (SRDS-17)*, pages 235–242, West Lafayette, IN, Oct. 1998.
- [13] A. Ranganathan, R. H. Campbell, A. Ravi, and A. Mahajan. Conchat: A context-aware chat program. *IEEE Pervasive Computing*, 1(3):51–57, 2002.
- [14] G.-C. Roman, Q. Huang, and A. Hazemi. On maintaining group membership data in ad hoc networks. *Technical Report wucs-00-26*, Apr. 2000.
- [15] M. Satyanarayanan. Pervasive computing: Vision and challenges. *IEEE Personal Communications*, pages 10–17, Aug. 2001.
- [16] A. Schmidt, M. Beigl, and H.-W. Gellersen. There is more to context than location. *Computers and Graphics*, 23(6):893–901, 1999.
- [17] M. Weiser. The computer of the twenty-first century. *Scientific American*, 265(3):66–75, Sept. 1991.
- [18] S. S. Yau, S. K. S. Gupta, F. Karim, S. I. Ahamed, Y. Wang, and B. Wang. Smart classroom: Enhancing collaborative learning using pervasive computing technology. In *ASEE 2003 Annual Conference and Exposition*, Nashville, Tennessee, June 2003.
- [19] S. S. Yau, F. Karim, Y. Wang, B. Wang, and S. K. Gupta. Reconfigurable context-sensitive middleware for pervasive computing. *IEEE Pervasive Computing*, 1(3):33–40, Sept. 2002.