massachusetts institute of technology — artificial intelligence laboratory

Activity Zones for Context-Aware Computing

Kimberle Koile, Konrad Tollmar, David Demirdjian, Howard Shrobe and Trevor Darrell

AI Memo 2003-015 June 2003

Abstract. Location is a primary cue in many context-aware computing systems, and is often represented as a global coordinate, room number, or Euclidean distance various landmarks. A user's concept of location, however, is often defined in terms of regions in which common activities occur. We show how to partition a space into such regions based on patterns of observed user location and motion. These regions, which we call activity zones, represent regions of similar user activity, and can be used to trigger application actions, retrieve information based on previous context, and present information to users. We suggest that context-aware applications can benefit from a location representation learned from observing users. We describe an implementation of our system and present two example applications whose behavior is controlled by users' entry, exit, and presence in the zones.

Sponsor: This research is funded by MIT Project Oxygen.

1 Introduction

The implicit control of applications using passively sensed context cues frees users from the burden of explicitly specifying details such as for whom a service is being performed, where it should occur, and how it should be delivered. Location is one of the primary context cues in context-aware systems, e.g. [6], Many of these systems define location in geometric terms such as a 3D (or 2D) position coordinate, room identifier, or a set of distances to known landmarks. While this definition is useful, it neglects how people actually use physical space.

One way to understand the use of physical space is in terms of zones that are created by elements of physical form and that support human activity, e.g., [3, 17, 21, 24, 25]. Architects, for example, design eating places, gathering places, and places of repose [3]. Anthropologists and sociologists talk of "hidden zones" in offices [17]. CSCW researchers have studied public and private zones for use of communication media in home environments [27]. In workspaces, we think of zones for quiet, solitary work; zones for informal meetings; zones for formal presentations [13, 33]. A table and accompanying chairs create a zone in which we might meet with colleagues. A whiteboard creates a zone in front of it; we draw on the whiteboard while standing in this "whiteboard zone".

We make two key observations about identifying such zones: Thinking bottom up, we can identify zones by observing people go about their daily activities, partitioning a space based on people's locations and motion. Thinking top down, we can define a taxonomy of zones based on prototypical human activity and physical form that supports that activity. Such zones, identified using either or both of these approaches, can be thought of as representing regions of similar context. With such activity-dependent zones, one can build more useful context-aware computing applications by (1) identifying meaningful zones to which users can attach semantics and preferred application behavior, e.g., behavior that would be triggered upon entry, exit or presence in a zone; and (2) enabling inference of human activity at various levels of abstraction, thereby improving the relevance of context-aware applications.

In this paper we describe our implementation of a system that identifies "activity zones" semi-automatically using a computer vision-based approach. We present a prototype system in which transitions between zones successfully control application behavior. Our system identified zones that were sufficient for inferring sets of similar activities, though not particular activities. We believe that knowledge of physical form (e.g., furniture type and location within a zone), combined with top-down knowledge of mappings between physical form and human activities, will enable inference of human activity at levels of abstraction higher than those enabled with an activity zone representation alone.

In the following sections, we review previous work, present the concept of activity zones, then describe our prototype system and ongoing and future research efforts.

2 Previous Work

The study of context and its role in ubiquitous computing systems is an active research field, and definitions for context and context-aware abound (e.g., see [11, 12] for surveys of definitions used in ubiquitous computing research; see [31] for survey of context in artificial intelligence research). Central to the notion of context is location, since many applications and services are conditioned on the place where they should be performed or displayed. Location-aware computing has become a topic of active research, and many schemes have been proposed for providing location cues to mobile devices using IR, RF, Ultrasound, and computer vision tracking systems. (For a survey, see [22].)

Early systems provided ad-hoc representations of location, usually tailored to specific device output. Recently a general scheme for device independent location representation and sensor fusion has been proposed, using a layered abstraction model [23]. The majority of location-awareness schemes report one or more of the following: raw 2D or 3D position information, room identity, or proximity to a beacon. These cues are useful for many tasks, but they are indifferent to the physical form or use of the space.

Room identity, raw position, or proximity are not a sufficient location representation in some cases. A person may be equally close to a whiteboard or a table, for example. If a display system knows that the person is standing and moving back and forth, it can infer that the person is at the whiteboard and choose a nearby wall as a display location, as opposed to the computer monitor on the table.

A few systems for location awareness are able to report information about sub-regions or furniture, and/or adapt over time based on observed behavior. The EasyLiving system allowed a map of regions to be used by the system to indicate places in the environment associated with specific context features, usually furniture [7]. These maps where manually drawn and provided to the system, and were not learned from observed behavior. Similarly the Sentient Computing System's notion of spatial containment [2] allows bounds for 2D position of active devices, e.g., an oval area in front of a computer display. A system for automatically mapping an environment based on the movement of personnel in a location-aware environment was described in [19]. This system was adaptive and learned from observing user behavior, but formed a map of a large scale environment and did not find contextually relevant sub-regions within rooms.

In our current research, we represent context by means of regions that we call activity zones, which are identified by observing users' locations and motions. A similar concept is discussed in [29]. A person's context is his or her entry, exit, or presence in one of these zones. In contrast to previous work, our representation is fine-grained (i.e., smaller than a room) and is learned from observing patterns of user behavior. It supports applications that can make productive use of inferences about a user's activity beyond simple position and proximity cues, without requiring the user to draw a map. We develop our scheme in an interactive framework: we learn the geometry of the regions from observed behavior, but rely on the user to associate semantics or rules with events related to those regions, as described in later sections.

We next describe the concept of activity zones, then discuss our implementation.

3 Activity Zones

Physical form—e.g. walls, furniture—partitions space into zones that are places of human activity. Walls create a zone that is a room; furniture creates a zone in which people might read or talk. In the floorplan shown below there is a zone created by the doorway, a zone created by the sofa, and a zone created by the corner desk, table, and chair.

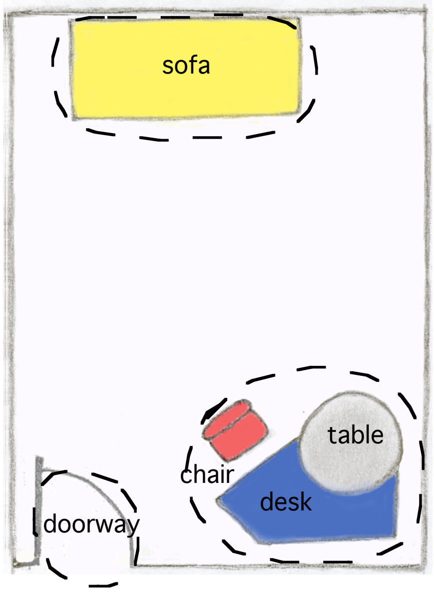

Fig. 1. Three zones created by physical form

Zones defined by physical form are useful as representations of context, but only partially capture a person's context—they ignore a person's use of a space. Rather than building a model of a space and its furniture as in [8], one can partition a space into zones by observing human activities. We call these partitions activity zones; e.g., the area in which a group of people are standing in a hall talking. If we take "observing human activity" to mean recording people's locations and motions, then for a floorplan such as the one shown in Figure 1, we would ideally find the zones shown in Figure 2: zone 1 is a region in which someone stands, zones 2 and 3 are regions in which someone sits, zone 4 is a region in which someone walks. Note that zones 1 through 3 correspond to the physical zones shown in Figure 1. They contain extra information, however—whether the person sits or stands in these zones. Zone 4 corresponds to an access corridor that represents a person walking between the three other zones. To identify such a zone in a model of the physical space, such as that shown in Figure 1, one would have to represent circulation paths, then make estimates about where people would travel along those paths.

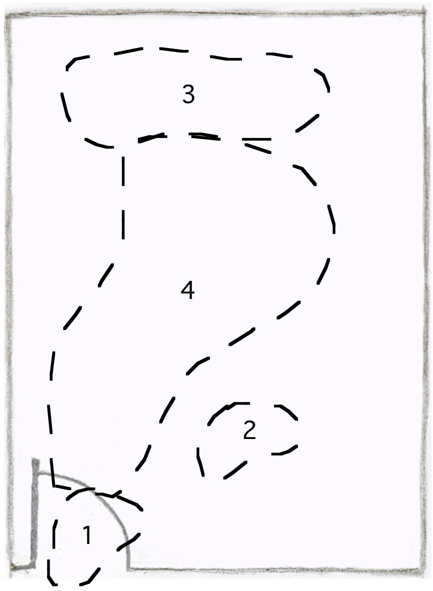

Fig. 2. Four zones identified by observing human activities

Both physical zones, shown in Figure 1, and activity zones, shown in Figure 2, can be thought of as representing regions of similar context. Physical zones represent location contexts; activity zones represent location and motion contexts. Physical zones could be inferred from observation of static furniture configuration, but activity zones need to be learned from statistics of user behavior. Application behaviors can be controlled by a person's entry, exit, or presence in either type of zone. Activity zones, with the extra information about user motion, enable application behaviors to be tied more closely to what people are doing rather than just where they are.

To use activity zones, a tracking system of some sort observes people's activities over time, say a week, and constructs a map of activity zones. A user (or eventually machine learning program) may then attach preferred application behavior to activity zones. A user also might attach semantics to the activity zones so that application behaviors could be specified for types of zones instead of individual zones. If a user, for example, did not want to receive phone calls while reading, she might label a particular zone as a reading zone, and indicate that while she is reading calls should be held. Once preferred behavior has been specified, the tracking system posts events about people's entry, exit, or presence in particular zones. An accompanying notification system informs interested applications of the events, and the applications react accordingly.

In this description of activity zones, the activity zone map is static—it was created by observing people in a space, then used assuming that the arrangement of space did not change. Yet zones are often correlated with furniture location, and in today's workspaces furniture is often moved [13, 33]. What happens, for example, when a sitting zone no longer contains the chair it once did? We map furniture locations to zone locations when an activity zone map is created, then periodically check that the current furniture locations match those in the current zone map. Our ongoing work on what we call dynamic activity zones is discussed in the Implementation section.

4 Scenarios

The following scenarios illustrate interaction made possible by the activity zones concept. In Section 5 we describe our implementation of two portions context-aware application behaviors mentioned here, including device control and the selection of display location and method based on context.

We consider scenarios taking place in a workspace containing zones such as those shown in Figures 3 and 4. Note that the zone near the round moveable table in Figure 4 is what we have called a dynamic activity zone.

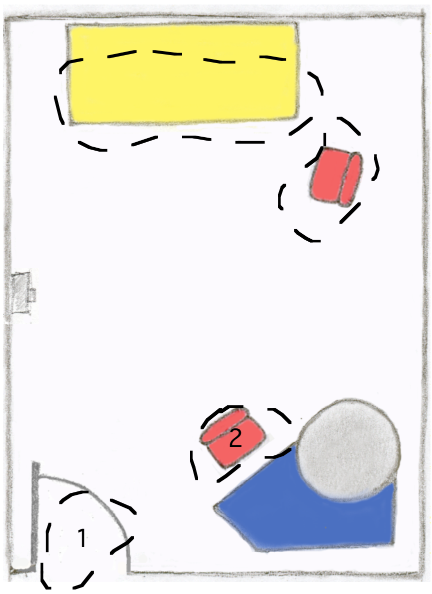

Fig. 3. Consider two zones, one at doorway one at corner desk chair

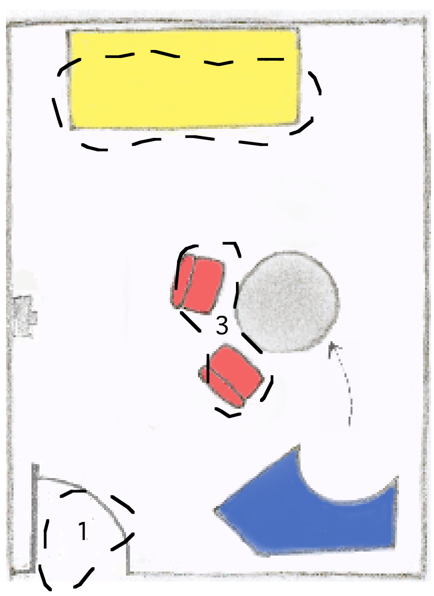

Fig. 4. New zone is created when round table is moved

Scenario A: Jane walks into her office and the overhead lights turn on. The room greets her and asks if she would like to be reminded of her calendar events for the day. She says yes, sits down at her desk, and her calendar is displayed on the computer screen near her. She notices that a student, Lauren, is coming by soon for a meeting about a term project. Jane asks the room to retrieve the notes from her last meeting with Lauren and to project them when she and Lauren start their meeting. She tells the room that she is finished with the calendar, and the calendar disappears.

Scenario B: Lauren arrives, Jane greets her and invites her to come in. Jane moves the table into the middle of the room and invites Lauren to sit down at the table. As Lauren sits down, the notes about her project are projected on the wall between the table and desk. Jane and Lauren start to discuss the project.

Scenario C: After discussing the project, Lauren leaves. As Jane moves the table back to its original location, the projector turns off. Back at her desk, Jane notices that the message light on her phone is lit and asks the room to play her phone messages for her.

In Scenario A, two preferred behaviors were triggered upon entry into the zone at the doorway: “illuminate the room” (or “turn on overhead lights”), and “ask about showing calendar”. Jane’s entry into the zone near the desk triggered display of the calendar on the display device appropriate for that zone (the computer screen on the desk). Scenario A also illustrates the creation of a dynamic event trigger: Jane requests that the room do something when a future event occurs, namely that it display notes when her meeting starts. Note that in order to identify a meeting, the room would need additional knowledge. It would need to know, for example, that meetings happen when more than one person is in zone 3, or meetings happen at tables and there is a table in zone 3. Without this additional knowledge, Jane would have had to request that the room display the information when she and Lauren entered zone 3.

Scenario B illustrates the display of information in an appropriate place based on the meeting starting in zone 3. Simple proximity to devices would not have worked in this example, because Jane and Lauren were equidistant between the projector and the computer display on the corner tables.

Scenario C illustrates that phone calls are held during meetings. It also shows the projector turning off when Jane moves the table back to its original location.

Together these scenarios illustrate the use of context to trigger room actions automatically, to retrieve information, and to present that information to users. (These uses of context are similar to those discussed in [8] and [12].) We describe our implementation of device control and an information display application in Section 6. We next describe the implementation of our system.

5 Implementation

We have implemented the activity zone concept as part of a larger system that provides services in an intelligent environment. Our activity zone system embodies a perceive-reason-act paradigm and is organized using a blackboard architecture [14]. Perceptual systems—the tracker and applications such as a message delivery system—post events to a blackboard. The blackboard provides a shared context memory and represents the current state of the world, e.g., the activity zone map in use; the people in the space and their contexts, i.e., location, motion (represented by height and velocity), and entry, exit or presence in particular zones; pending and processed events. A reasoning component does forward inference, abstracting events into higher level context statements, e.g., a meeting in zone 3, or mapping events to requests for action, e.g., turn on the lights. The requests for action are sent to device controllers and other applications, which in turn process the requests.

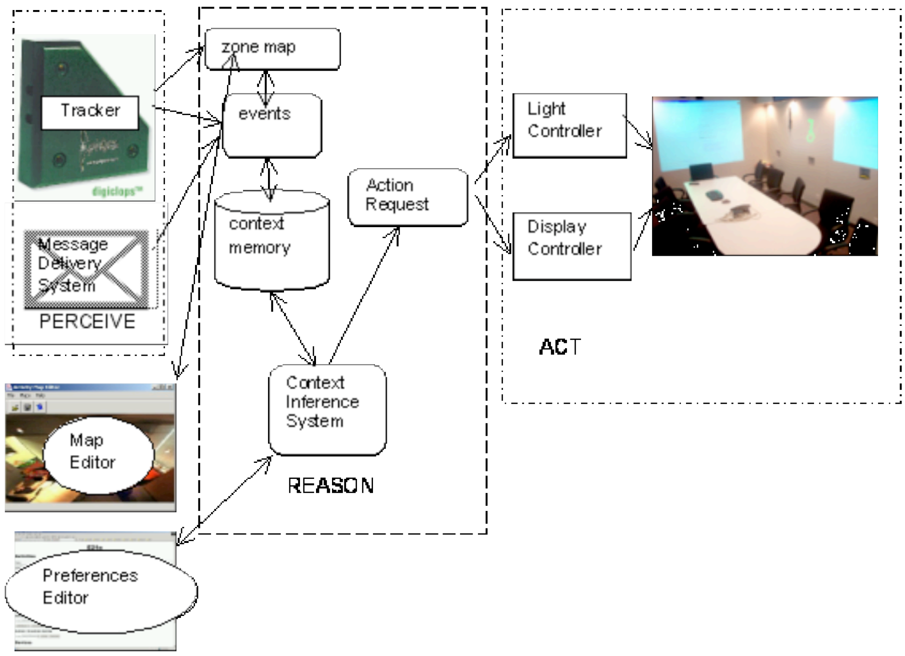

Fig. 5. System architecture

Figure 6 is a sketch of the floorplan we adopted for the workspace used in testing our system.. Figure 7 is a photograph of the space. The furniture can be moved around easily to create zones for individual work and collaborative work.

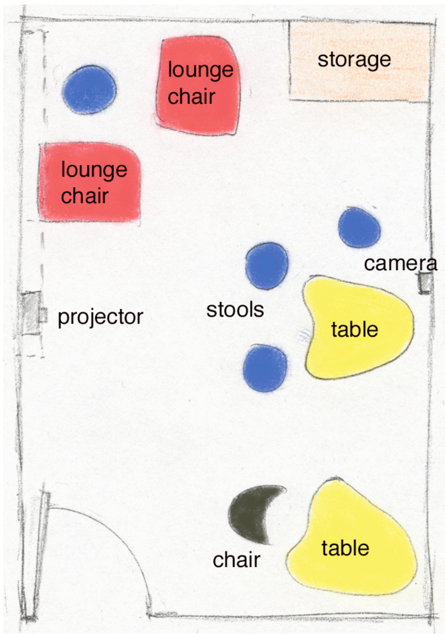

Fig. 6. Floorplan sketch of workspace

Fig. 7. View of workspace from doorway

5.1 Tracking System

To support activity zones, a person tracking system must provide information in real-time on the number of people in a space and each person's location and motion (e.g., represented by height and velocity). These are general requirements, and many different tracking approaches can be used, including those that track wearable or handheld devices using IR, RF, infrared, GPS, 802.11 [22]. We have chosen to use a passive person tracking system based on computer vision, since it does not require any additional devices to be worn or carried by users. The concept of activity zones and our techniques for constructing and manipulating, however, are independent of the tracking system used, as long as it provides the information specified above.

We use a multi-camera stereo-based tracking system to track people in indoor environments based on the work described in [10]. This person tracker provides a history of 3D information for every person in the observed space. The information is a triple (x, y, h), where x, y are the coordinates of the person in the ground plane, and h is the relative height of the person with respect to the floor. Since tracking data are time-stamped, the instantaneous velocity (v_x, v_y, v_h) can be derived. We then characterize a person at location (x, y) using the activity feature $f(x,y) = (h, v, v_{lt})$, where h is the height, v is the instant ground plane velocity norm, and v_{lt} is the average ground plane velocity norm over a certain period of time. By using the feature $f(x, y)$, we can capture the configuration (sitting, standing) and movement of a person over both short and long periods of time.

To estimate an activity zone map, we track people in a space for a period of time, collecting a dense set of activity features $f_i(x, y)$ and location features (x, y), then segment the activity features using a two-step clustering process. We first cluster the activity features into classes, each representing a similar activity (i.e., configuration and movement), which is represented as an average activity feature F_k. Then for each class, we cluster the associated (x, y) features into regions. The resulting regions represent activity zones, regions in 2D space that are characterized by an average activity F_k. As different activities may happen at the same location, activity zones may overlap. Once an activity zone map has been created, a person's entry, exit, or presence in an activity zone is identified by matching the person's instantaneous location and activity features to the map.

To allow heterogeneous agents to manipulate and access our activity map data structure, we represent it using simple XML primitives. The figure below shows an example of a portion of an activity map for our workspace.

```
<?xml version="1.0"?>
<amap xsize=200 ysize=200 im="835-PTGO.jpg">

<zone id=0 label="desk" height=1.1 velocity=0.1 color="ff0000" >
                12,182 12,184 38,188 …
</zone>

<zone id=1 label="table" height=1.7 velocity=0.5 color="00ff00" >
100,182 100,180 89,154 …
</zone>

<zone id=2 label="lounge" height=0.8 velocity=0.05 color="0000ff" >
150,130 150,132 135,118 …
</zone>

</amap>
```

Fig. 8. Activity-map in XML. Labels are user supplied

The length of time for collecting data for an activity zone map is an open research question; we typically track people for a day. People can be detected with an accuracy of about 20 to 30 centimeters. This parameter is what limits the performance—if people are sufficiently spaced, the tracker can track all of them.

Figure 9 shows an example of the data clustering for an activity zone map.

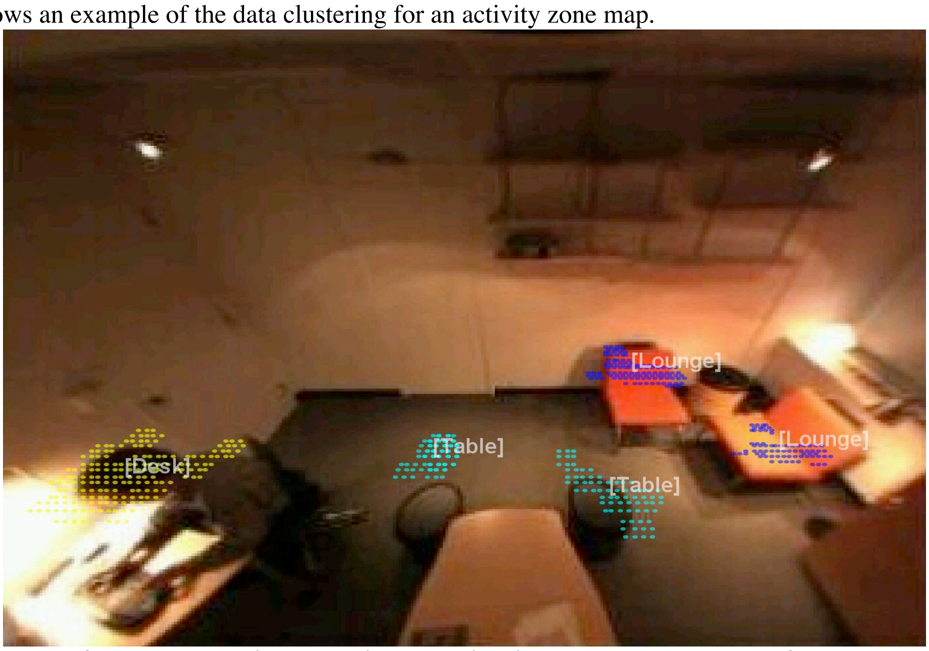

Fig. 9. The workspace from the camera's perspective, overlaid with clustered tracker data for three zones; labels are user-supplied

The above figure shows three activity zones clustered around the desk, meeting table, and lounge chairs. The meeting table zone is the union of smaller zones that the tracker identified around the stools. (The issue of small zones is discussed further in the Example Applications section.) The tracker found a total of 11 zones, which represent four functionally separate areas. The four areas are the three shown above, plus an access corridor through the middle of the room (To make the figure more easily readable, the access zone is not shown.)

We noted earlier that an activity zone map may not be relevant if a physical space is rearranged. In such circumstances, how does the system either swap in an alternate activity zone map or know to compute a new one? We have two answers: (1) save an image of a space for each of several activity zone maps, and periodically check that the image matches the current furniture configuration, or (2) record furniture type and location in or near particular zones, then use a furniture tracker to notice a mismatch between original furniture locations and new locations. We have implemented suggestion 1; we plan to try suggestion 2 next. Figure 10 shows clustering in an alternate activity zone map.

Fig. 10. The workspace from the camera's perspective, overlaid with clustered tracker data for 3 zones; zone labels are user-supplied

5.2 Blackboard and Infrastructure

A blackboard provides a shared memory through which system components communicate. It contains a context memory, a set of zone maps, a current zone map, a perceptual event queue, and a requested action queue. We use an implementation similar to that described in [32], though in our current prototype the handling of events and action requests is outside the blackboard.[1] Perceptual systems, such as the tracker, post events to the blac k-board. A context inference system then reasons forward from the events, using such rules as "if there are more than 2 people in a zone then there is a meeting in the zone". Resulting inferences, such as "person 1 in a meeting," are stored as statements in the context memory. Representing context at this higher level of abstraction allows users to specify preferences more naturally, as mentioned in the previous section's scenarios. It also enables preferences to work in the presence of dynamic activity zones since preferences can be associated with types of zones, e.g., meeting zones, rather than particular zones anchored in physical space. The preferences are entered using a simple web-based interface.

In addition to posting context statements, the context inference system may post requests for action on an action request queue. These requests are sent to device controllers and other interested applications by means of an agent-based infrastructure, e.g., [9, 16, 18, 26, 30]. An agent-based infrastructure provides a mechanism for managing software components, including device controllers, enabling communication between the components; and protecting security and privacy. We currently are using its publish-and-subscribe messaging facility as a communication layer to control devices and applications in the workspace. The facility provides a means for software agents, which represent sensors and services, to register interest in particular classes of notifica-

[1] Our events and action requests are not processed sequentially, as the term "queue" implies; each runs in its own process. The blackboard control structure is a focus of our current research.

tions. An agent that controls a particular light in a particular activity zone, for example, might listen for notification of a "turn on light 1" event. Light controller agents currently send "turn on," "turn off," "dim," and "brighten" notifications to agents controlling particular lights.

6 Example Applications

We describe two implemented examples of using activity zones as context cues. The first example, control of lights and audio uses a direct mapping from events to environment actions. The second example, delivering a message based on context, demonstrates the use of inference to hypothesize higher level context descriptions.

6.1 Device control

The system adjusts light and audio levels when people enter and exit particular activity zones. We used our tracking system to generate activity zones as described in the previous section. When we examined the resulting activity zone map, we discovered that many small zones were in the same functional region, e.g. each stool was in its own zone. We found it more useful to think of the stools as a single zone, especially when it came to specifying preferred device behavior. We implemented a zone editor that allowed us to group small zones into larger ones. We selected a subset of zones, one of which was the larger one created for the table and stool area, then added labels and preferred light and audio settings for each zone. These settings were specified using a rule-like language and interface provided by the agent-based infrastructure. Typical device settings set by a user were:

- When someone enters desk zone, turn on the lamp and music
- When someone enters lounge chair zone, turn on the lamp in that zone and turn off the lamp in desk zone
- When a second person enters table and stools zone, turn on the projector and turn off the music

Tracker events for entry to and exit from zones were then mapped to requests for device control by means of the preferred settings. Figure 10 shows the state of lights and projector when two people are sitting in one of the zones.

Fig. 9. Two people in the meeting zone

We informally observed people working in the space. We noted that the activity zones did correlate well with particular activities; typing at the keyboard in the desk zone, reading and talking with others in the lounge chair zone, working alone or with others at the table in the meeting zone. Informally and qualitatively, we found that people who used the room liked the layout of the room and the feedback provided by the tracker display of clustered data. Most people thought the automatic control of lights, radio, and projector novel and useful. People described the zones by either the objects in the zone or by the activities that took place in the zones. Rather than having the user always add labels for zones, object recognition and activity inference can be used to provide labels and to insure that labels carry semantic information rather than being just symbols. The following example applications sow categorical inference with activity zones.

6.2 Message delivery

We implemented a message delivery system for context-aware notification to test the usefulness of activity zones and higher level inference. Knowledge about what constitutes a useful zone can be specified in the context inference system described above. If we assume, for example, that we care about a zone C that is the union of zones A and B, then the inference system could have a rule such as "if a person enters zone A or zone B then the person enters zone C." When the tracker posts an event to the blackboard that a person has entered zone A, the inference system sees that event and infers that the person has entered zone C. If there are preferred application behaviors attached to entry in zone C, the inference system then can post the appropriate requests for action.

Examples of context inference and message delivery rules adopted by us are shown in Figure 10.

```
if number of people in zone ?z > 2
then meeting-in zone ?z

if person ?p in zone ?z
   and meeting in zone ?z
then person ?p in meeting

if person ?p in meeting
   and message-delivery-event for ?p ?msg
then deliver-without-display ?p ?msg

if deliver-without-display ?p ?msg
   and person ?p in zone ?z
then notify agent for default-display-device for ?p in ?z
            "deliver-without-display" ?msg
```

Fig. 10. Example context inference and message delivery rules (in pseudo-code); ?x indicates a variable

We had two people sit on the stools at the table, and then sent a message to one of them. [2]The tracker noticed a person in each of two stool zones; the context inference system inferred that each person was in the meeting zone and that there was a meeting in that zone. The message delivery system then displayed an icon on the computer screen in the desk zone, which was the default display device for the recipient when she was in a meeting. With one person in the room, the message delivery system displayed the message on the computer screen if the person was in the desk zone (where the screen was), and it displayed the message using the projector if the person was in the lounge chair zone. Quantitative user studies of the comparative usability of this system are in progress. Anecdotally, experiences of novice users (computer science graduate students) suggest that the system is a natural and effective way to adjust notification state. Users need not make explicit gestures or utterances to change notification state; they need only specify once the preferred behavior of the environment.

The above description illustrates two important points: (1) activity zones can be used to deliver context-dependent information, and (2) the blackboard and context inference system enable delivery preferences to be stated in terms of high level context descriptions.

7 Discussion and Future Work

The concept of activity zones—regions of location context formed by observing user behavior—is a broad one, and we have only begun to explore its full extent. In addition to the ongoing user studies of the current prototype mentioned above, we anticipate being able to use information about furniture identity and locations to augment our context inference system with simple activity models representing such information as "meetings often happen at tables," "informal meetings often happen in comfortable chairs." We could use the perceptual features from our tracking system, plus identity of furniture objects, to index into a catalog of higher level contexts representing such activities as being in a meeting or reading (e.g., as in [31]). With extra information about objects and activities, we would be able to infer users' activities more accurately, thus increasing the relevance of task-related information and services provided by the workspace.

In building a catalog of higher level contexts, we plan to investigate activity theory [5, 15], research on understanding how people work [13, 28] and research on how workspace design affects how people work [13, 20,

[2] People were already identified by the system with an explicit utterance.

33]. This body of literature will provide us with examples of the kinds of social interaction that help define context for people.

We are keenly aware that many people are uncomfortable having their activities observed by cameras, and privacy issues deserve attention when designing context-aware systems (e.g., [1]). In our work to date we have found that it is important to have computationally transparent perceptual systems so that users are aware of the kind of image information that is processed and stored. Abstract trajectory data (with voluntary and explicit identity assertion) is generally considered less invasive of privacy than full motion video. We plan to investigate methods for real-time visualization of tracker state.

We have shown that our concept of activity zones can provide valuable context cues for context-aware applications. So far we have explored this idea with relatively simple applications—device control and message delivery. We plan to continue our investigations by using activity zones with other applications, in particular retrieving information based on a previous context. We also plan to explore the idea's utility in retrieving information based on a predicted context. A zone's information about motion, for example, could be used to predict where people are likely to be sitting or standing. This predictive capability may prove valuable with tasks that are more structured than those usually found in an office environment. In a laboratory setting, for example, scientists often engage in an ordered set of activities, centered around various pieces of laboratory equipment. Context-aware computing researchers are building applications to support such laboratory activities, e.g., [4]. The activity zones in such a setting would center on laboratory equipment stationd. The zones could be mapped to an ordered list of activities, which then could be used to predict likely transitions between zones. By anticipating what scientists may do next, a support system could initiate start up procedures when necessary or gather additional information needed at a particular station. In these and other applications, we believe that activity-based location regions provide a richer level of context information than has been previously available in context-aware computing.

References

1. Ackerman, M., Darrell, T., and Weitzner, D.J.: Privacy in Context. Human-Computer Interaction 16 (2001) 167-176.
2. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., and Hopper, A.: Implementing a Sentient Computing System. IEEE Computer 34 (2001) 50-56.
3. Alexander, C., Ishikawa, S., and Silverstein, M.: A Pattern Language: Towns, Buildings, and Construction. New York: Oxford University Press (1977).
4. Arnstein, L.F., Sigurdsson, S., and Franza, R.: Ubiquitous computing in the biology laboratory. Journal of the Association for Laboratory Automation (JALA) 6 (2001).
5. Bannon, L.J. and Bødker, S.: Beyond the Interface: Encountering Artifacts in Use. In: Designing Interaction: Psychology at the Human-Computer Interface. J.M. Carroll (ed.). Cambridge University Press, New York (1991) 227-253.
6. Beigl, M. (ed. Special Issue on Location Modeling in Ubiquitous Computing. Personal Ubiquitous Computing. Vol. 6 (2002).
7. Brumitt, B., Krumm, J., Meyers, B.B., and Shafer, S.: Ubiquitous Computing and the Role of Geometry. IEEE Personal Communications 7 (2000) 41–43.
8. Brumitt, B., Meyers, B., Krumm, J., Kern, A., and Shafer, S.: EasyLiving: Technologies for Intelligent Environments. In: Proceedings of the 2nd International Conference on Handheld and Ubiquitous Computing: Springer-Verlag (2000).
9. Coen, M., Phillips, B., Warshawsky, N., Weisman, L., Peters, S., and Finin, P.: Meeting the Computational Needs of Intelligent Environments: The Metaglue Environment. In: Proceedings of the 1st International Workshop on Managing Interactions in Smart Environments (MANSE 99). Dublin, Ireland: Springer-Verlag (December 1999).
10. Demirdjian, D., Tollmar, K., Koile, K., Checka, N., and Darrell, T.: Activity Maps for Location-Aware Computing. In: Proceedings of IEEE Workshop on Applications of Computer Vision (WACD2002) (December 2002).
11. Dey, A.K. and Abowd, G.D.: Towards a Better Understanding of Context and Context-Awareness. In: Proceedings of the What, Who, Where, When, and How of Context-Awareness Workshop,CHI 2000: ACM (2000).
12. Dey, A.K., Abowd, G.D., and Salber, D.: A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. Human-Computer Interaction 16 (2001) 97-166.
13. Duffy, F. and Powell, K.: The New Office. London: Conran Octopus Ltd (1997).
14. Engelmore, R. and Morgan, T. (eds): Blackboard Systems. Addison-Wesley, Reading, M A (1988).

15. Engestrom, Y., Miettinen, R., and Punamaki, R.-L. (eds): Perspectives on Activity Theory. Cambridge University Press, New York (1998).
16. Grimm, R., Davis, J., Lemar, E., MacBeth, A., Swanson, S., Anderson, T., Bershad, B., Borriello, G., Gribble, S., and Weatherall, D.: Programming for Pervasive Computing Environments. University of Washington, Computer Science and Engineering: Seattle, WA (2001).
17. Hall, E.T.: The Hidden Dimension. New York: Doubleday (1969).
18. Hanssens, N., Kulkarni, A., Tuchida, R., and Horton, T.: Building Agent-Based Intelligent Workspaces. In: Proceedings of The International Workshop on Agents for Business Automation. Las Vegas, NV (June 2002).
19. Harle, R.K. and Hopper, A.: Using Personnel Movements for Indoor Autonomous Environment Discovery. In: First IEEE International Conference on Pervasive Computing and Communications (PERCOM 2003). Fort Worth, TX: IEEE (2003).
20. Hartkopf, V., Loftness, V., Aziz, A., Shankavaram, J., and Lee, S.: The GSA Adaptable Workplace Laboratory. Lecture Notes in Computer Science 1670 (1999) 29-44.
21. Hertzberger, H.: Lessons for Students of Architecture. Rotterdam: Uitgeverij Publishers (1991).
22. Hightower, J. and Borriello, G.: A Survey and Taxonomy of Location Systems for Ubiquitous Computing. University of Washington, Computer Science and Engineering: Seattle, WA (2001).
23. Hightower, J., Brumitt, B., and Borriello, G.: The Location Stack: A Layered Model for Location in Ubiquitous Computing. In: Fourth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 2002). Callicon, NY (2002).
24. Hildebrand, G.: The Wright Space: Pattern and Meaning in Frank Lloyd Wright's Houses. Seattle: University of Washington Press (1991).
25. Hillier, B. and Hanson, J.: The Social Logic of Space. New York: Cambridge University Press (1984).
26. Johanson, B., Fox, A., Hanrahan, P., and Winograd, T.: The Event Heap: A Coordination Infrastructure for Interactive Workspaces. In: Fourth IEEE Workshop on Mobile Computing Systems & Applications (WMCSA 2002). Callicoon, NY (2002).
27. Junestrand, S., Keijer, U., and Tollmar, K.: Design Patterns for Private and Public Digital Spaces in the Home Environment. International Journal of Human-Computer Studies 54 (2001) 753-778.
28. Kirsh, D.: The Context of Work. Human-Computer Interaction 16 (2001) 305-322.
29. Lahlou, S.: Observing Cognitive Work in Offices. In: Cooperative Buildings: Integrating Information, Organizations, and Architecture. N. Streitz, et al. Editors. Springer, Heidelberg (1999) 150-163.
30. Martin, D.L., Cheyer, A.J., and Moran, D.B.: The Open Agent Architecture: A Framework for Building Distributed Software Systems. Applied Artificial Intelligence 13 (1999) 91-128.
31. Turner, R.M.: Context-Mediated Behavior for Intelligent Agents. Int. J. Human-Computer Studies 48 (1998) 307-330.
32. Winograd, T.: Architectures for Context. Human-Computer Interaction 16 (2001) 401-419.
33. Zelinsky, M.: New Workplaces for New Workstyles. New York: McGraw-Hill (1998).