

Multimodal and Ubiquitous Computing Systems: Supporting Independent-Living Older Users

Mark Perry, Alan Dowdall, Lorna Lines, and Kate Hone

Abstract—We document the rationale and design of a multimodal interface to a pervasive/ubiquitous computing system that supports independent living by older people in their own homes. The Millennium Home system involves fitting a resident's home with sensors—these sensors can be used to trigger sequences of interaction with the resident to warn them about dangerous events, or to check if they need external help. We draw lessons from the design process and conclude the paper with implications for the design of multimodal interfaces to ubiquitous systems developed for the elderly and in healthcare, as well as for more general ubiquitous computing applications.

Index Terms—Interaction design, multimodal, older people, pervasive computing, seniors, ubiquitous computing.

I. INTRODUCTION

IN THIS paper, we describe some of the lessons that we have learned from the design of a ubiquitous multimodal system that supports an independent home life for older people. While several “intelligent homes” and architectures to support eldercare have been documented, many of these have focused on the support of leisure activities, home automation, and support for communication (e.g., [1]–[3]). Very few systems that support the home life and healthcare of older people have been developed to improve quality of life and the alleviation of risk, with a few exceptions (see [4]–[6]), although even these have often been designed as “add-ons” to generic “Home of the Future” automation systems, rather than systems primarily designed to support independent living for older people.

There are a number of reasons why users might find having different forms of input and output (I/O) useful when interacting with computer systems. Computer systems have moved a long way from purely keyboard–screen interactions to include a range of other interactional techniques. Our aim here, though, is not to discuss the forms of multimodal I/O devices or the relative merits of each, but to explore *how* the notion of multimodality affects interaction between the user and the computer system. In an attempt to highlight these issues, we will discuss a computer-based system—the Millennium Home—in which interaction with the system can be conducted through a variety of different media, depending on the user's physical or contextual requirements.

Manuscript received September 15, 2003; revised May 11, 2004 and June 29, 2004. This work was supported in part by EPSRC Grant GR/R09503/01.

M. Perry, L. Lines, and K. Hone are with DISC, Brunel University, Uxbridge, Middlesex UB8 3PH, U.K. (e-mail: mark.perry@brunel.ac.uk).

A. Dowdall was with Brunel University, Uxbridge, Middlesex UB8 3PH, U.K. He is now at 4A Hampden Square, Southgate, London N14 5JR, U.K. (e-mail: alan.dowdall@hotmail.com).

Digital Object Identifier 10.1109/TITB.2004.835533

The rationale of the system was to provide a safety net for independent-living older people. It is important to note here that the system, as described, is intended to support users that are not currently cognitively or physically impaired, but who may become ill or injured in the course of their home life and will require some form of assistance. If a problem situation were detected through embedded sensors, the system would attempt to contact the resident to warn them about the situation and to establish whether they required assistance or not. The system was intended to provide appropriate, context-sensitive information to the resident where it could be sensed and acted upon without the need for locating, or carrying, I/O devices for interaction. Resident–home interactions and dialogues were designed using appropriate interactional *modes* for the user's location and activity. This involved the design of a multimodal interface that made use of visual, speech–audio, appliance, and environmental activity to drive I/O actions. While this kind of system supports a relatively simple (but nontrivial) problem in a multimodal interface design (assisting older people in living independently), it raises a number of issues, and has many of the problems that other ubiquitous systems with a wider scope of application might have.

II. HEALTHCARE, SMART HOMES, AND USER INTERACTION

A. Lifestyle Monitoring and Older People

Sensor-based “lifestyle,” or health, monitoring has a huge potential in supporting an independent life for older people for a variety of quality of life, economic, and social reasons. The value of these kinds of systems is made clear when it is seen that currently only 10% of the population of “older people” live in accommodation with a form of institutional support [7]—when a much larger proportion of the elderly population could potentially benefit from some form of technologically mediated care or support. Assistive healthcare technologies (and in particular, lifestyle monitoring) can, therefore, offer older people the ability to experience the benefits of sheltered housing from their own home, giving them some control and flexibility over their private living arrangements.

The problem of providing effective healthcare for older people is one that is also likely to exacerbate over time. The proportion of older people in the population is increasing in developed countries such as the U.S., Japan, and Britain because of the demographic structures of these countries, and this has led to problems in supporting the quality of life in this aging population. In Britain, healthcare improvements and lower birth rates have resulted in those aged 65 years and older, representing 16% of the population [8], and this is expected to

rise to around 19% in the first quarter of the this century. For reasons of personal choice and economics, many of these older adults will continue to live independently, with a large number of them living alone. Even though support for those living independently is cheaper than in residential care, healthcare and welfare organizations are still being financially stretched to provide an acceptable level of community support for these people.¹ Clearly, technological support that would continue to help older people live independently would be of enormous financial and, perhaps more importantly, social benefit.

Elderly people are particularly prone to accidents and falls in the home and can often lie injured and undiscovered for long periods of time. On a somewhat different level, they tend to perceive themselves as likely victims of crime, leading to a low quality of life resulting from this fear of crime [9]. Cognitive problems can also occur, one of the most typical reported in the literature being an increased level of forgetfulness (see, for example, [10]), which can cause healthcare problems as people forget to take regular medications. These are critical issues to support in the design of augmentative systems. However, in addition to their increased levels of physical frailty and cognitive problems, older users may also face various other problems, such as hearing and motor limitations, which may impact on their ability to interact with a computer system. Of course, there is huge variation in the skills and abilities of older people, and these may even vary over time for a single person, so that this user group poses an unusual and interesting set of design questions for technology developers.

B. Sensors and Interaction in Smart Homes

Porteus and Brownsell [11] describe the development of healthcare monitoring systems as falling into three generations, each with an increasing level of sophistication in terms of their sensory functionality and levels of user and community interactivity. First-generation systems include technologies such as community, or “pendant,” alarms, which the user must actively press to initiate a call for assistance. These systems have a number of limitations, not least that the user may be unable or unwilling to initiate the alarm. Second-generation systems attempt to improve on the limitations of the first-generation systems through their application of a level of “local intelligence,” such as the use of sensors to proactively detect alert situations. They then go on to describe a future scenario involving third-generation systems that can provide access to commercial services, “virtual neighborhoods,” and communities of support for the residents of these homes. While they do not explicitly mention sophisticated content-appropriate interactivity as part of the definition of a third-generation system, an interactionally simple, yet content- and context-relevant interface, will be required in order to make these interactive services possible or practically usable.

The Millennium Home system documented in this paper falls into this third-generation category of telecare, building on work

¹Allied to this is also an increasing body of legislation on social services for this age group: This means that they have to provide an acceptable “duty of care” in monitoring the health and ensuring the safety of these residents. In itself, this has implications for the design of systems to support older people, which may have to respond to yet another stakeholder group.

from the Anchor Trust/BT Telecare project [11], [12]. The Anchor Trust/BT Telecare project is identified as a second-generation telecare system, collecting data on the resident of the home to build up a “normal” profile of the older person/resident’s daily activity patterns, for example, when the resident would go to bed, use the refrigerator, or any unusual fluctuations in the house temperature (see also [13] for a similar system). Wide deviations from this profile would trigger an alarm alert resulting in a very simple, automated telephone message to the resident, in which they would be made aware of the nature of the alert, and cancel it or request help from a carer (with no reply to this also triggering an alert to the carer). However, while user feedback to the system in the evaluations was positive, the system could be over-sensitive, and to a degree, intrusive in its alerts (many of which were false alarms). The problem of providing a quick and low effort mechanism to cancel false alarms, and conversely (but as importantly), or providing a quick response to a real emergency (rather than waiting for a specified “time out” before raising an external alarm to a carer) should be seen as at the center of such an enterprise. Indeed, the role of the interface in any such system is seen by Porteus and Brownsell as crucial:

“The design and operation of the user interface is perhaps the primary way of ensuring that the user keeps control of the technology and their surroundings. Speech and vision are the most natural forms of communication and are, therefore, obvious targets for a user interface . . . Whatever form of interface is chosen, it must be suited to the needs of the particular individual for which it is intended, and the social and psychological aspects associated with its use must not be overlooked.” [11, p. 56].

This critical issue of user interaction forms the starting point of the design of the Millennium Homes project, developing a sophisticated and sensitive approach to user interaction with the computer system monitoring the user’s activities and controlling the alarms. The selection of appropriate forms (or modes) of communication with the user, and usable interactive dialogues for each of these modes of communication form the crux of this work. By giving users the ability to make use of different modes of communication (known as multimodal interaction), we can develop “natural” and context-sensitive interaction, yet this raises important questions of which modes are appropriate (and when they are appropriate) and of how to generate interaction dialogues for these different modes that are useful and meaningful to users. All of this is complicated by the particular needs of the elderly users of the system, in that they may have a variety of cognitive and physical problems, and that they may also have suffered some form of accident that could constrain their interactive abilities with the system.

C. Mode-Switching and the Adaptable User Interface

This section examines the issue of developing multimodal systems that allow an appropriate mode of user interaction. One of the key advertised benefits of multimodal systems, across a range of application domains, is that of “robustness.” It is argued that multimodal systems offer the opportunity of increasing the accuracy of interaction in otherwise error prone activity settings and interactional contexts. This is possible because users can

proactively hone the media that they use in their interactions to make the best use of the modalities of those media, or reactively move to another medium if required (for example, see [14]–[16]). Indeed, there is much evidence to suggest that combinations of complementary technology can be used to increase recognition rates (e.g., [17]–[19]), although these have mainly been examined with reference to speech and direct manipulation–gesture.

While there is evident value in the use of multimodal interaction to support recognition robustness through the “complementarity” that is advocated by these researchers, this is only a part of the more general value that multimodality can bring to the utility of interactive systems. One of the added values that interactive multimodal systems can bring is *mode-switching*. Mode-switching allows the users or the computer to change to another mode of interaction if the current mode is (or has been recognized as being) a problem. Mode-switching differs from multimodal complementarity (noted earlier) in that it involves the use of a single interactional modality at a time rather than the simultaneous use of different modalities. This allows the selection of an interaction style and medium that is appropriate to the current setting or activity, and which works around its constraints (such as its acceptability or environmental noise). A simple example of such contextual issues occurs when driving a car or walking, when screen- and mouse-based interactions may be inappropriate. Decisions that users make on the choice of an appropriate media for interaction matches the ways that we interact with other people, honing our responses and providing cues to our conversants to elicit a particular response from them. Such systems that support mode-switching are described as having “adaptable user interfaces” [20]. While there are few documented instances of adaptable user interfaces being deployed outside of the desktop environment, it is useful to define the term here: An “adaptable user interface is defined as an interface that:

- 1) supports a number of different dialogue modes. . . ;
- 2) allows the user to switch between dialogue modes at any time, i.e., even in the middle of a command;
- 3) makes the switch between dialogue modes smoothly and naturally;
- 4) makes it easy for the user to learn how to use the different dialogue modes. . .” [13, p. 1353].

As can be seen from the final two points, it is clear that even in applications that just run on a personal computer desktop, making smooth transitions between modes and learning how to use the different modes is considered to be important. These are perhaps *the* critical usability issues to address in the interaction design of mode-switching systems.

This issue of mode-switching in multimodal interaction is a particularly important one for the design of ubiquitous or pervasive (we use the two terms interchangeably here) computing. With microsensitive location monitoring and contextually aware systems that can sense a person’s location and their current activity (as in a healthcare system such as the Millennium Home), we can design both the medium of interactivity and interaction dialogues to “fit” the user’s requirements and activity more ap-

propriately. This is the approach that we have chosen to use in our designs.

One of the crucial issues in designing a mode-switched multimodal system is determining how the selected mode of user interaction maps onto the *methods* that users can employ to achieve their intended system outcomes (i.e., their goals). Selection of an appropriate mode is important in user performance and satisfaction. For example, Grasso *et al.* [17] examined task completion time, error reduction and system acceptance in multimodal systems, showing that improvements could be made if system responses were matched to users’ perceptual motor skills. To put it another way, the mode of interaction with the computer has a clear effect on its usability. Different modes of interaction have different properties and potentialities. For example, navigation through verbally presented menu structures is very different from navigation through visual menus due to effects of the signal on the cognitive system. The reason for this is that auditory cues are short-lived, but are good for making users aware of state changes to the system (i.e., they afford reminding *in* time); conversely, visual cues are useful in that a system’s state is always available and accessible (i.e., they afford reminding *over* time), but state changes may be missed if a user’s attention is elsewhere.

While these differences may appear to be obvious, there is an important and often ignored concomitant design-related issue that arises as a result of this: Interactions with different media may require different dialogue structures (i.e., the *methods* of the paragraph above) to achieve the same outcomes. To design a usable mode-switching multimodal system, simply transforming the interactional medium is not (always) appropriate—the content of the interaction may need to be adapted to fit the constraints of that modality. However, to do this without losing consistency in the interface design is problematic, and this is recognized as one of the critical factors in human–computer interaction and usability design (see, for example, [21]). Use of different interaction methods, even when carried out in the same medium to achieve the same outcome, can be confusing to users, as they may have learned one method of interaction and have expectations of system behavior that are not met through other methods of interaction.

Another significant problem that arises out of mode-switching occurs when there is a need for changing the mode of an interaction *during* an ongoing interactional sequence. When the methods of interaction for the different modes do not neatly map onto one another, how can the system adjust its dialogue to meaningfully continue its interaction with the user, without restarting the interactional sequence again (and the consequent confusion to the user that this might cause)?

When mode-switching is determined by the system (and not by the user), there can also be problems in determining the choice and appropriateness of the interactional modality for any given situation. The computer system needs to be able to determine the most appropriate modality for the interaction, while allowing users to select or change the modality of interaction to suit their own needs or requirements. The contextual appropriateness of an interactive medium is of *specific* relevance to the Millennium Home, where users may be unable to use a particular interactive medium or method for reasons of accident or

illness (i.e., the very reasons for using the system), as well as less health-critical reasons including the speed at which older users can interact with the system (e.g., to initiate or cancel an alarm), and their mobility/comfort in using it, as they may not feel like jumping up to answer the telephone, or having to move to a visual display in order to make an input onto it. Mode switching, therefore, offers a powerful tool for enabling pervasive healthcare for older users; however, implementation of such a mode-switching system could potentially present users with additional interactional complexity, and in designing the system, we need to ensure that this overhead is low, and at a minimum that is not greater than the advantages that these systems could bring to their users. This is a major concern driving the design of the interface to the Millennium Home.

D. A Caveat: Premises and Limitations

While the Millennium Home system presented here has been developed in accordance with our own user studies (some of the results of which are presented in [9]) and the existing research literature on elders and eldercare, we recognize that the model of the older user that has been used in our interaction design is based on a relatively broad characterization of the user population. As the reviewers of the paper have pointed out, like other large groups of users, older people have a wide range of behaviors and idiosyncrasies that may impact on the acceptability of the Millennium Home system, the ways that it is used by them, and its effectiveness in supporting their needs. However, in order to develop a system that is acceptable to older users with a range of interactional, behavioral, and healthcare-related issues, a relatively simple user model has been used for the purposes of this paper. Nonetheless, determining the needs of the older users of healthcare systems such as the Millennium Home is an important area of research, and one that continues to drive thinking on future instantiations of the system. Indeed, as the Millennium Home system develops and we are able to get feedback in evaluations from greater numbers of real users, we will be better able to determine how effective this model is in meeting the needs of a diverse set of users in naturalistic settings and whether more sophisticated user models may need to be developed. However, this does not detract from the interaction techniques and mechanisms described here, and these are likely to require relatively minor tweaks and extensions to the interaction design rather than extensive redesign. The conceptual foundations underlying the interaction architecture and the interactional techniques used, and the contextual issues and implications for design presented, are likely to continue to remain valid and relevant to any new developments to the Millennium Home, and by extension, to similar ubiquitous systems.

III. TECHNICAL AND INTERACTIONAL DESIGN

This section of the paper describes our experience of developing the interface to a third-generation telecare system—the Millennium Home. Our intention is not to provide details of the technical structure of the system, other than to demonstrate where these have an impact on the interaction design; rather, the coverage here is to discuss the rationale and design of the system as it will appear to its users—through their interaction

with it. The key issues covered in this section encompass the implementation issues surrounding the development of pervasive multimodal systems, and in particular our approach to the selection of an appropriate mode of interaction. It then continues with a discussion of the interactional issues that arose from the technical architecture of the system, with specific emphasis on the development of mode-specific dialogue and interaction.

A. Project Goals, Constraints, and Implementation

In order to implement a lifestyle monitoring system that supports some of the needs of independent-living older users, the Millennium Home project partners (please see acknowledgment section) have developed a sensor-based system that allows the detection of a resident's activity within their home. This includes sensors that detect movement, doors and windows being opened and closed, falls, temperature, and a range of other events within the home.

In the design of the resident-home interaction, we were given an explicit design requirement that the system should not require the resident to carry anything, such as a pendant alarm, health monitoring devices, or personal location identifiers. Along with the novel research demands that this places on the design, we have qualitative evidence that leads us to believe that older users often will not carry such devices on their person. One reason given by existing pendant alarm users for not carrying their devices was that they were worried that they might break the devices if they fell over.² There is also a danger that falling onto a carried device could do the user more harm than if they fell directly onto the ground. This design constraint places a heavy demand on the system to respond to events and to do this at an appropriate time.

As information about the resident's activity in the house is collected through its network of environmental sensors, it is processed, and when judged appropriate, an interaction with the resident is initiated either via audio or visually, depending on their activity and/or location. This allows the resident to act directly on their environment to rectify the situation (e.g., shutting an open door), to cancel the alarm sequence using their voice (through a voice recognition interface), or by interaction using an input device (through their television set or over a telephone keypad). If there is no response, or an inappropriate response is detected to an alarm, an external alarm message will be raised to an external call center so that help can be sent to the resident.

For the purposes of the design, we have scoped out several problematic and complex environmental settings, so that these additional complications and design constraints will not impact on the prototype system. In a sense, most of these are reasonable and practical design constraints to place on a final design, as we (and any future implementers of such a system) are likely to be able to specify the kind of houses and some of the domestic configurations that the system can be used within. The system is, therefore, designed for a small house with its own front door and no alternative entrance. The system is also not designed to be used in homes with large pets, because these might be incorrectly sensed by the systems as the resident and an inappropriate system response generated. Finally, the technology

²Personal communication, Prof. Heinz Wolff, Brunel University, U.K.

was not intended to be a diagnostic tool or a “health” monitoring system, nor was it intended to be a method of domestic technology management. These could be envisaged as add-on aspects of the system, but are explicitly not addressed in our designs. Indeed, people with serious health-related problems or cognitive and physical disabilities are likely to have particular and specialized requirements of the system. We have, therefore, designed the system to support a cognitively fit and able-bodied user, excluding the very old and those with dementia, because of the additional interactional demands this would place on the user interface. This meets the requirements of a home that will monitor critical health-critical events and support simple and appropriate interaction with older people that are healthy and able-bodied within the norms of their age group.

We have discussed user requirement definitions and the prototyping process elsewhere [9], and do not cover them in this paper as they are not relevant to the question of multimodality in a ubiquitous environment. The interaction design of the system in terms of its interaction structure and dialogue, and the issues surrounding the choice and implementation of this is, however, of relevance here, and these are discussed below.

B. Technical Infrastructure

In order for the system to assess user activity so that it can raise an alarm or select an appropriate mode of interaction, sensors monitor the “state” of the resident and the house. These sensors include the following:

- 1) passive infrared sensors to detect movement (including specialized motion sensors that can detect if the resident falls over);
- 2) pressure sensors underneath the legs of chairs and beds to detect whether the user is on them;
- 3) burglar alarm style sensors on windows and front door to detect if they are in open or closed state.

In addition, the system is fitted with adjustable timers to allow medication alerts to be raised at an appropriate interval, and temperature sensors to check that the home is not getting too warm or cold for the resident’s health. However, it is important not to focus on the sensors in themselves, as the interaction design is not linked directly to the sensors, but the interpretation of the sensors (several of which may need to be read together to be meaningfully interpreted)—and this is managed by a separate system. The outputs of this interpretation are fed into the Millennium Home’s interface module, and this is the component that we discuss in the rest of the paper. A final point that needs to be made about the sensors is that their state-based information is dynamic—as new sensor readings of state changes are detected, this affects the nature of the interaction.

In the design of the human–computer interface module to the Millennium Home system, communication from the system to the resident can take place through a variety of mechanisms. These include the following:

- 1) a computer-activated telephone, which can ring the resident (and is linked to the Millennium Home interface through a digital exchange);
- 2) loudspeakers (placed in all rooms);
- 3) television/monitor screen (computer activated).

In turn, the resident can communicate with the system through the following:

- 1) telephone, via an audio menu (with speech recognition) or using a key-press menu operated system;
- 2) screen-based interactions through menu selections over a television/monitor screen using a remote-control device;
- 3) activation of environmental sensors (e.g., shutting a door or window, or by the user standing up);
- 4) limited voice recognition. Situated microphones can detect vocalizations, although this is limited by poor recognition accuracy; however, noisy domestic equipment can be automatically turned off through the X10 protocol.³ Voice recognition is, therefore, only used in situations where other methods are cumbersome or otherwise problematic—currently this is only used in the bathroom.

As previously noted, the range and variety of the different media involved has led to a number of interesting issues in designs that link interaction between them (e.g., management of menu navigation across different audio–visual media).

C. Interactive Dialogue Structures and Sequences

Following on from the system requirements, dialogues with the resident are rarely a matter of simple yes/no answers, but will often involve extended interactions. The nature of these dialogues will depend on the complexity of the required interaction, the location of interactive events within the home, the resident’s visual, verbal, or hearing limitations, and the speed with which the resident can act on any of the system’s requests.

This section details the dialogue sequences. Rather than go through a huge range of possible system interactions, we will initially illustrate this with a simplified system initiated dialogue (Fig. 1). In this case, the Millennium Home system (identified in Fig. 1 as “SYS”) has detected an alarm state, and places a telephone call to the resident. Throughout the sequence, the dialogue allows three levels of functionality: 1) to repeat the system output; 2) to cancel the alarm state; 3) to request human assistance. Please note that the dialogue structures shown below just document the semantic content of the interaction (i.e., their meaning, but not their actual wording) and the user options that will achieve the functionality required (i.e., call for assistance, cancel the alarm, or repeat the question).

Although a number of I/O modalities are available, for the purposes of the illustration here, the example dialogue sequence presented uses the telephone as the sole mode of I/O. This set of dialogues is generic and simplified, but does illustrate the kind of dialogues that users can have with the system. However, the example does not demonstrate any of the complexity of the system in detail to show how multimodal interaction could be conducted.

The diagram in Fig. 2 illustrates some of the functionality and complexity involved in a more complex set of fleshed out dialogues in a medication alert scenario using the telephone, screen–screen controller, or loudspeaker–control device as media through which to interact. The medication alert scenario may appear complex, but it is one of the simpler scenarios

³[Online]. Available: <http://www.x10.org/>

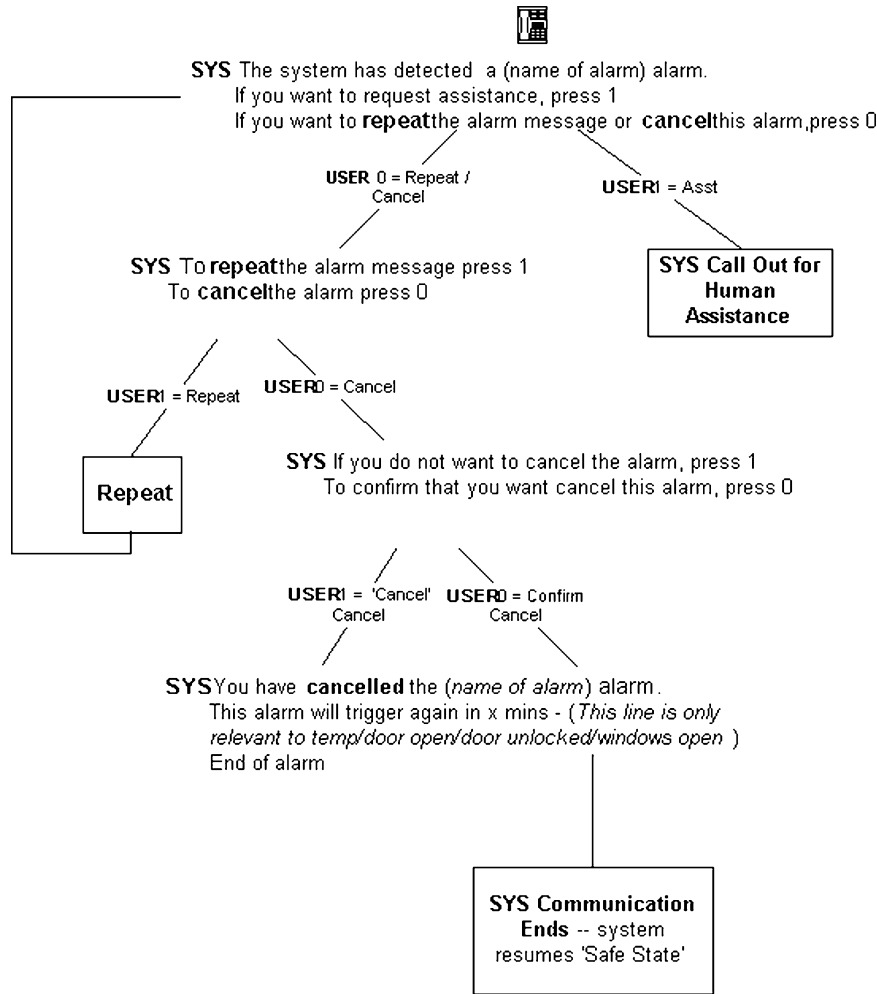


Fig. 1. Simple system dialogue structure.

developed—others involve a great deal more interaction and complex interdependencies.

While it is not necessary to fully understand the design formalisms used in Fig. 2 that represent the interaction sequences or the specifics of the design, it is useful to examine the similarities and differences between the structure of the top three shaded boxes that represent different modes of interaction, and the interaction pathways and sequencing between system events. Note that while the different modes of interaction are functionally equivalent, the *methods* that are used to control and manipulate the state of the system differ from one another because of the *qualities* of the different media used in response to the alarm context.

Working through the interaction in Fig. 2, the flowchart represents the possible states and state transitions following the start of the medication alarm. First, a countdown is initiated: This time will be used to assess whether to raise an external alert if the alarm is not dealt with quickly enough by the resident. An appropriate mode of communication is then determined (assessed through a variety of factors, such as the resident’s location or their current activity), and communication then takes place with the resident over the loudspeaker (users can interact through pressing buttons on the telephone—for archaic reasons, labeled as the “IPB”), over the telephone, or through their tele-

vision set (interacting with it through their remote control). At various points, other alarm sequences may be activated or returned from, and when an ambiguous response is detected; these processes are not represented fully in the diagram (these hidden sequences are labeled as X, X1, and A). Throughout the process, the countdown is polled to check on the time spent since alarm initiation. If no activity is recorded on the loudspeaker or television, the system defaults to the telephone. If there is no response before the time-out, an external alarm is raised, and feedback is given to the resident that help is on its way (“post communication processes”).

In themselves, the sequences represented in Fig. 2 are not particularly interesting to the topic of the paper—what is relevant here is the differences in interaction between the modes, and how the selection of a mode and movement between these modes is managed in the interaction.

D. Interaction Design

In developing the interaction sequences⁴ discussed in the previous section, the design presented us with several interesting

⁴Here, we distinguish between interaction sequencing and interaction dialogues; an interaction sequence relates to the *flow* of events in an interaction as the user navigates through the interaction, while an interaction dialogue is the *content* of an interaction as presented to the user (as text or speech).

questions about the important factors that needed to be considered. These factors are likely to be highly relevant in the design of multimodal interfaces to ubiquitous computer systems in other contexts. Some of these are discussed below, using the Millennium Home as a context.

Initiation modality: the primary mode of communication employed by the system. This is dependent on the resident's activity (as detected) in the Millennium Homes system. Thus, if the resident were watching television, this would be an appropriate initiation modality. In time- and health-critical situations, broadcast speech is used to initiate communication with the resident over loudspeakers. Audio-speech is a quick method of alerting the resident to the fact that the system is aware of the alarm state and it may also be the only plausible mode of communication, given that the user may be unable to move. If the resident does not provide feedback (i.e., an input response that the system can understand) then the system should continue to use speech throughout the interaction.

Dialogue location: where dialogues with the system physically take place. In the Millennium Home, dialogues take place in the location where the alert signal was raised, or if resident movement was detected in a particular location, at that place. When no response is detected from the resident, subsequent dialogues are broadcast throughout the residence. However, for other ubiquitous computing systems, general broadcasts may or may not be appropriate.

User-input modality: the mode of communication to be employed by the user. This can be free-format, in which the user makes a decision, or constrained through the system's requests for a particular user response (for example, contact that is initiated over a telephone and requires a telephone-based response). In the Millennium Home, only speech and deliberate physical actions (such as the press of a button or close of a window) will return the system to a safe state—clearly indicating that the resident is conscious and able to respond appropriately.

Dialogue completion: provision of feedback following interaction. After completion of a set of dialogues with the computer system, users need to be clearly made aware that the dialogue sequence has been completed and is not in a mid-cycle pause. In the Millennium Home, the system's state is continuously made available through the content of the system's dialogue. The system also provides reassurance to the resident, informing them that a carer has been notified after a call has been made for outside assistance. In the Millennium Home, residents are able to check on current and past interaction events by picking up the telephone receiver or accessing the system log over a dedicated monitor.

Intradialogue temporal sequencing: these are the system pauses between acts of communication within a dialogue sequence. In the Millennium Home, where dialogues are conducted using an input device, residents are given a 3–5-min period with which to locate the device before attempting subsequent communications or contacting external parties. This time may vary depending on the urgency of the alarm, and the duration of the ongoing alarm situation.

Interdialogue action sequencing: in systems where multiple interaction events can happen simultaneously, communication

of this information to the user needs to be segmented into meaningful interactive sequences. Such sequencing is not necessarily best dealt with through a temporal ordering of events, as some events may have more importance or urgency than others. In the Millennium Home system, conflicts between alarms can arise, and these are addressed through an assessment of alarm importance with respect to the resident's safety. The system needs to be able to deal with multiple simultaneous alarms, without creating a confusing mass of unrelated, but temporally sequential, dialogues. For example, when a window alarm with a 30-min dialogue warning cycle is interrupted by a fall alarm with a 5-min dialogue warning cycle, precedence is taken by the more urgent alarm—the window alarm is turned off until the fall alarm has been dealt with.

Interevent dependencies: interaction events may have dependencies on one another—changes to one may affect the others. For example, in the Millennium Home, incoming newly raised alarms can change or terminate ongoing communications with the user. This follows the previous point, in which some simultaneous situations of concern may be interrelated, for example, a lack of movement and a subsequent fall. Thus, we may have simultaneous alarm interactions that require modification to standard user interaction sequences or dialogues. Our simple solution to this has been to address this through an urgency monitoring system that records the frequency of previous recent alarms and any concurrent alarm events, and subsequently increments (or decrements) the level of urgency (see below). This information is used to moderate the time period between dialogues.

Urgency management: allowing the system to adapt itself through accessing historic context. In the Millennium Home systems, there is a requirement for the system to address problems as quickly as possible, while at the same time, minimizing the number of false alarms. With increasingly failing health or additive health problems, problems can escalate quickly, and some kind of urgency management may be required to expedite an alarm that under “normal” circumstances would otherwise not be urgent. One solution to this is to use the number and types of alarms previously generated as a potential indicator of the instability of a resident's health, and to increase the urgency of dealing with these accordingly. According to the system's assessment of danger to the resident, we have, therefore, designed our system to moderate the times between repeated dialogues within interactive sequences. Although not implemented, more complex situation-sensitive solutions could be used to dynamically modify the mode of dialogue presentation, the number of warning messages, and the content of the dialogues according to the previous alarms generated.

IV. CONTEXTUAL ISSUES

Six key high-level contextual issues in interaction design emerged from our involvement in the design process. While several of these have been directly addressed in the previous section, these issues are interesting in that they clearly relate to the design of ubiquitous systems across a broad range of problem domains. Please note that not all of these have been addressed in our own design, although through our design and

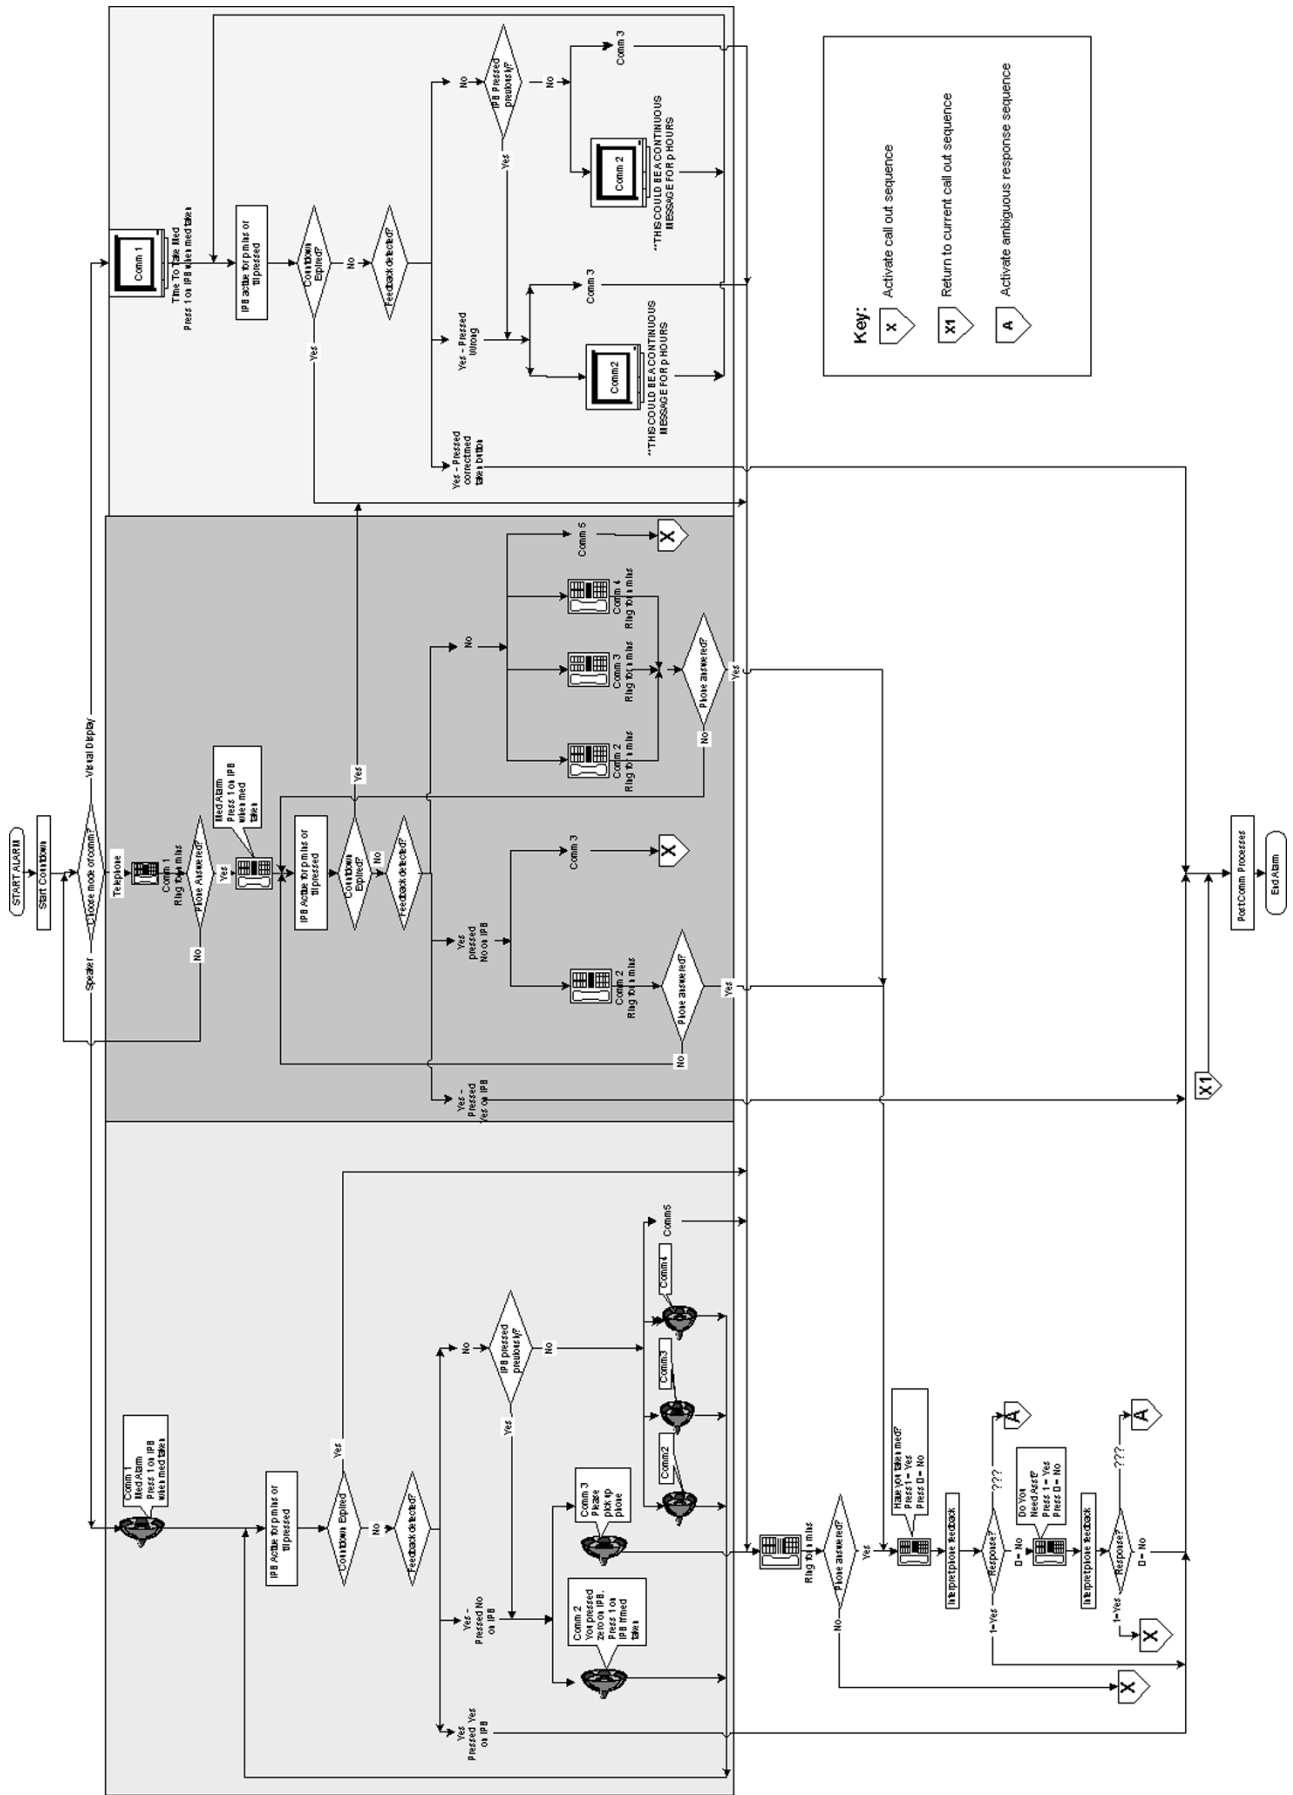


Fig. 2. Medication reminder alarm flowchart.

user evaluation processes, they have emerged as highly important.

- 1) *Location-appropriate interaction*: Interaction needs to occur in an appropriate location and form for users to respond to. This may require attention to be paid to the dialogue content and mode of interaction to make them context-appropriate. The Millennium Home takes information from the sensors to determine *where* the resident is, for example, in the kitchen, bathroom, etc., and this is reflected in the system's dialogues and mode of interaction. Thus, if the user is sitting down with the television turned on but has not moved for an unusually long time, the ideal modality would be to conduct an interaction using an overlay on the television screen in the room that the resident was in, and not to broadcast this over audio to every room in the house.
- 2) *Activity-appropriate interaction*: Similarly to the above issue, where possible, interaction sequences and modes should be honed to the user's current activity. The Millennium Home interprets information from the sensors to assess *what* the resident is doing, for example, cooking, watching television, on the toilet, etc. In the example of the resident sitting down and watching television without moving for a long time, alarm interactions need to be initiated in such a way so as not to significantly disrupt their enjoyment of a television program.
- 3) *Multiuser interaction*: Computer systems may be used by more than one user, and this can greatly complicate interaction. Multiuser systems present something of a problem for many ubiquitous computing technologies: It is not always obvious for a computer system to detect *which* user has issued a command or requires a system output. It is possible to give each user a unique identifier (perhaps some form of radio transmitter embedded in a watch or pendant), but it is not always possible to ensure that they carry it with them. In the Millennium Home, we have attempted to determine when visitors arrive, switching the system off in these circumstances. However, in future designs, we may need to allow for interaction with multiple occupants. This would require the system to act differently to when the resident was alone. Although multiple occupancy of the home may be difficult to determine automatically, it is possible to determine who is present using a variety of activity pattern sensing techniques (see [22]), but it is also possible to determine multiple occupancy through dialogue with the user (although these dialogues may be annoying and intrusive if frequent or if this significantly adds to dialogue completion time). Another simpler solution is to make the system "fail-safe" (i.e., to assume single occupancy), although this may generate an increased level of false negative alarms (requiring the user to manually cancel inappropriate alarms). Awareness of, and support for, multiple occupancy may also improve user satisfaction and acceptance of the system. Within this design frame, we need also to be aware that not all passers-by of a ubiquitous system are permitted users. In the example of the Millennium Home, a burglar-intruder

presents a very different form of multiuser interaction and we need to ensure that they are recognized as such and cannot take inappropriate control of the system.

- 4) *Interaction over time*: Action develops over time, and there is usually a historic context to activity. This is certainly true in the case of systems such as the Millennium Home in which previous activity within an interaction cycle, and activities over several interactive cycles, may change the user's interaction requirements. As already noted, in the Millennium Home, a simple urgency management process is used, which adapts the system to gross changes in the number and frequency of alarms in the recent past.
- 5) *Priority assessment for interaction*: At any one time, a number of interaction events may occur concurrently that require the user's attention, and the system needs to have a method of assessing the order of priority for events that should be addressed in the foreground. Again, this has been addressed in the Millennium Home system (albeit in a simplistic way) through a hierarchy of alarm importance. Multimodal systems are useful in this respect, in that different media are available to alert users to different events—in most cases, an auditory medium will be the most appropriate medium for alerting users to the prioritised event.
- 6) *System-state availability*: So that users can adapt their own activities around a system's expected actions, users need to be aware of the current internal state of the system (or at least the part of that system relevant to the user). This is considered good design practice for interaction design in general. However, provision of information on system state can be difficult to achieve in embedded and ubiquitous systems, as the issue of "where" the interface is to the user is not always clear. To a limited extent, information on the system state is available in the Millennium Home system through its announcements. Thus, dialogues are designed to give feedback on past actions, current possible interactions, and future outcomes following an interaction; for example, system announcements follow a dialogue format such as this (note: this is not a actual dialogue):

< you have not responded to two warnings;
 answer yes or no to whether you require assistance;
 if you do not answer within n min,
 a call for outside assistance will be made > .

Information on the system state can also be directly accessed through the telephone interface and through a computer monitor, although this aspect of the design is not well advanced in the current version of the Millennium Home.

V. IMPLICATIONS FOR DESIGN

Clearly, the six key contextual issues mentioned in Section IV the previous section are central to the design of multimodal interfaces for ubiquitous computer systems, with implications

that go far beyond simply controlling an alarm. If we are to work with, and live within, environments fitted with sensors and allowing us to control appliances and other computer-based systems, where we are, what we are doing, who we are doing it with, what we have been doing, and what we judge to be important in—and what we can expect to be the outcome of—the interaction are clearly critical components that cannot be ignored in designing approaches to interaction. There are a huge variety of methods for collecting this information. Some of this information is readily available while others are less easy to obtain, but integrating them together and making use of this information to modify interactions is likely to be hard to do. Following experiences with our own design processes, we, therefore, present six implications for designing multimodal and ubiquitous systems.

- 1) When initiating contact with the user, the selection of an appropriate modality is important in making them aware of the need for dialogue, while not being unnecessarily intrusive. There is a difficult balance here between the system's perceived importance of the interaction, and the user's perception of the appropriateness of that level of intrusiveness for the event. In terms of the design, this problem can be ameliorated if the system has knowledge (even if in a simple form) of the potential importance of the dialogue to the user and the potential intrusiveness of the form of the dialogue to the user for the user's ongoing context. Another solution to this problem is the ability to escalate the intrusiveness of the interaction alert (for example, through mode-switching) in successive attempts.
- 2) The location that the dialogue takes place in is important in the contextual appropriateness of the interaction. Developing Point 1 above, it makes sense to sequentially increment coverage of the computer's initiation dialogue through an escalating level of personal to public visibility. However, for different kinds of interactions (and possibly, different kinds of contexts), the dialogue may (if important) begin in an escalated, or public state, or (if relatively unimportant or of a private nature), only escalate to a within a less publicly visible state. As before, dialogue systems need some knowledge of the importance of the interaction to the user, and the level of public visibility acceptable for the interaction.
- 3) User control is important in all forms of user interface design, but it is particularly difficult to design for ubiquitous systems, because of the lack of a unique control artifact (or even a visible control artifact at all; see [23] for a good review). Being able to cut short an interaction and knowing when the dialogue sequence is finished is important in allowing the user to know when they can fully disengage from their interaction. However, in systems where the technology is embedded within the world, it is hard for the user to expedite an interaction, or to know where or how they can obtain information on the system's current state. Information on system state needs to be easily available to the user so that they can move on to another activity without the worry that they are still "halfway" through an incomplete interaction sequence. This is clearly of critical importance in passkey-protected activities and in interactions that initiate other follow-on events.
- 4) In large or complex systems, there may be a need for users to be able to engage in several ongoing interactions at any one time with one or more networked devices. Some of these operations and devices will have dependencies between them so that carrying out interactions as separate operations may appear to be nonsensical to the user, for example, reminding a user to lock their front door, while at the same time informing them that the door has been knocked down by a burglar. This means that there may be a need for dialogues that can be modified according to these interactions and dependencies. One of the simpler techniques to deal with this may be through developing a system of priority management, allowing the system to present users with the issue understood to be of most pressing concern to them. Of course, this approach is not always appropriate, and design decisions will have to be made to allow the system to deal with these interevent dependencies appropriately.
- 5) Consistency needs to be considered across modalities, not just in the simple form of dialogue syntax. There needs to be consistency at the level of semantics, of meaning, in what the systems is attempting to perform through the interface. By this, we mean that it should be possible for the user to understand the current state of the interface (even if they have not encountered the current interaction sequence through this modality) based on their knowledge of the same interaction through other modalities. In our interface, for example, this is achieved by giving the user three warnings prior to alarm callout with a similar set of system requests for user interaction for all of the interaction modalities—this can be seen in the similar interaction structures of the three modes shown in Fig. 2.
- 6) Modal switching within an ongoing interactive sequence is problematic for interaction designers, because different media have different cognitive properties. So, an interaction initiated aurally with verbal user feedback that later shifts to a screen and button-based mode (either when directed by the user or in reaction to events) will place different demands on the user. Interaction sequences that are optimized for one mode will not necessarily be appropriate for the other (hence, the difference in the interaction pathways for the three modes seen in Fig. 2). One solution to this is to segment interactions into small units, so that each smaller interactional unit can be completed (or fail), allowing following interactions to be conducted in any other modality. As can be seen from our own designs in Fig. 2, the system allows mode-switching at points of interaction failure with the user. Small, self-contained interaction sequences, therefore, allow greater flexibility in the mode of interaction.

A final issue for discussion lies in the development of the system for users with impairments. Although the system was not originally designed for use by users with cognitive impairments (such as users with Alzheimers, or following a stroke) or physical disabilities (impaired mobility or wheelchair bound), a pervasive healthcare system, such as the Millennium Home, which makes use of mode switching in interaction does have the potential to offer these users some benefits. While this is more

speculative than the material covered in the rest of the paper, this might be achieved by two main techniques, 1) the provision of support for specific problems faced by impaired users and 2) designing the system so that the mode selected is appropriate not just for the general context, but for the nature of the impairment (e.g., by developing new interactive modalities) so that it is tuned to the abilities of the users. Focusing on these in more detail.

- 1) There may be alarm types specific to particular groups of users with disabilities, from those with physical disabilities (e.g., visually impaired users who need to be warned about obstructions), to those with cognitive problems (e.g., strokes that have resulted in impairment to their language comprehension or production, and requiring different forms of alert warning or input methods). These may require different sensors to be embedded in the home and linked to the system, and will almost certainly require different interaction sequencing and forms of dialogue (such as grunt-based acknowledgment). Many of the same contextual issues and design guidelines mentioned in the implications earlier may be relevant in developing the interaction designs for this, but additional factors may impact on this and will need to be accounted for in a useful and usable interface design.
- 2) More appropriate user-appropriate mode selection could be achieved by changing the rules by which the selection of the mode of communication was made, to attune this to the particular needs and abilities of the user. For example, at a simple level, for users with visual impairments, interactions requiring screen-based interaction could be removed from the mode selections; at a more complex level, for users with early onset Alzheimers, this could be manifested in the need to repeat the alerts more often, and simultaneously, over a range of different modalities (visual, audio, even possibly haptic) to allow for their increased likelihood of forgetting the alarm itself or the exact nature of the alert situation.

These areas of potential value for users with disabilities and their associated design suggestions are necessarily speculative, as we have no experience in developing or working with systems to support this user group, and the reader is cautioned in this respect. Nevertheless, this is an area that is clearly deserving of more research. In addition, detailed evaluations of any such technology developed would be required in order to assess the relevance and safety of augmentative systems for users with disabilities, especially given the very diverse nature of disability experienced by older people, both in terms of their severity and effect.

VI. CONCLUSION

At the core of the design of the Millennium Home system has been the central issue laid out by Porteus and Brownsell [11], to develop an interface to a healthcare system that allows a resident of the Millennium Home to control the technology and their environment. The system allows its older users—with

their particular interactional and healthcare requirements—to monitor and effect changes to the ongoing state of the system. This is achieved through interactionally simple dialogues over a range of modalities, giving them an increased level of control over their own healthcare provision, and allowing them to integrate the technology into the fabric of their own individual demands, routines, and needs. By giving older users the opportunity to make use of the most “natural” and appropriate modes of communication with the system for their current contexts of use, we have provided a powerful tool with which they can control an important component of their own healthcare provision. Yet beyond this healthcare application scenario, the paper has raised a number of issues that are pertinent to the development and application of interaction techniques for more general ubiquitous and pervasive computing systems, and these are also an important outcome of the work documented here.

Over the course of the paper, we have presented a relatively simple multimodal system in a particular instance of a ubiquitous computing environment. However, despite the constraints that this highly restricted domain brings to the problem, the design challenges have been hugely complex. A truly ubiquitous system (covering multiple user scenarios, a range of user ages and abilities, and supporting many tasks and activities) would have a far greater design space to cover, and it is hard to imagine how this could be done without building some form of intelligence into the system’s interaction. We have not attempted to do this for a variety of reasons (primarily, this not being necessary for our simplified scenario of use). However, building intelligence into systems introduces its own set of design problems, for example, how would an intelligent interface to a system present a consistent pattern of interaction if its interactional content and structure is generated dynamically? How would it address competing access to resources: dynamically, according to need, or formally, according to a predefined rule structure (and moreover, if the system applied dynamic rules, how would the users know about resource allocation)? Rather than present answers to these here, we leave this to other researchers. However, wherever consistency expectations are violated, it is vital that users are provided with good feedback to ensure that they understand what is happening in the interaction.

Throughout the design, we have attempted to avoid using the “computer” model of interaction, relying on the metaphors and patterns of interaction familiar to computer users, e.g., screen-based dialogues with terms such as “OK” and “Cancel” or goal-based interactions. Such interaction techniques are inappropriate, as users of healthcare technologies may not be experienced with computer systems. It needs to be remembered that users of the Millennium Home system will be older, with possibly less experience of information technology, from different social and cultural backgrounds, and with a varying level of education. This also holds true for a variety of ubiquitous systems that provide applications “off the desktop” [24], rather than through a traditional screen-based user interface. In these circumstances, methods of interaction may be more effective when based on other familiar forms of control metaphor, such as the telephone exchange, mechanical interactions, or forms of interaction based around exploration (such as games or search) rather than goal-based interaction (see, for example, [25]).

During the design process, we have attempted to put the particular concerns of older users at the center of the interaction design process. This has, for example, informed the nature of the repeated warnings in the system, giving users enough time to physically get up and engage with the selected mode of interaction (focusing on older users' physical abilities) and in repeating the interaction multiple times, thus, not forcing them to have to remember the exact nature of the alert (supporting older users' cognitive abilities). However, this research also raises a number of important issues that arise out of provision of ubiquitous–pervasive computing in other healthcare systems. Some of these are not covered in the body of this paper itself, but will arise out of the implementation of Millennium Home-type technologies in the homes of older people. These cover a gamut of concerns from privacy (who is going to be able to access this information and for what reasons?), security (can inappropriate people, such as burglars, deactivate the house alarm system?), and healthcare legislation (by providing healthcare equipment, might the provider, for example the local council, be legislated against if the system fails?). We can but be aware of these concerns in developing and implementing such systems in the real world, recognizing that the system described is currently a research vehicle implemented within a laboratory environment. Evaluations from the implementation of systems such as the Millennium Home in the “real world” outside of the computer laboratory are likely to provide us with new and possibly unexpected data regarding their use and perceptions about their utility, from all of the stakeholder groups, including residents, their carers, healthcare professionals, system managers, and so on.

Other issues concerning ubiquitous–pervasive computing in healthcare are directly addressed in the paper itself: Patterns of activity are not stable and are liable to change over time. Systems that cannot cope with fluctuations in their user's health are unlikely to be able to provide an appropriate set of responses. Issues of modality, the central aspect of work covered in the paper, are also important considerations in thinking about how the system will interact with its users—a single mode of human–computer interaction will be limited in its contextual appropriateness, or its ability to cope with a user's failing physical or mental capacities. An adaptable user interface—managed through the system's ability to perform dynamic mode-switching in interaction—is one way of achieving a suitable form of multimodal interaction design for this user group. The provision of feedback about the current state of the system is vital in safety-critical healthcare systems—and this is particularly hard to provide in ubiquitous systems because of the pervasive (i.e., embedded and often invisible) character of the technology. Feedback about event closure is particularly important, so that users are aware that they have completed everything that they need to and can relax (or conversely, need to continue with the interaction); in the case of the Millennium Home, users are made aware that either an alarm has been cancelled, or help called for.

The lessons learned from the design of this multimodal interaction system have much to offer the designers of other ubiquitous computing environments about developing useful and usable systems. It is possible that designing ubiquitous systems for novel environments, such as this, will help support the de-

sign of more general ubiquitous systems, because they present unusual scenarios that force us to think of new ways to achieve solutions for interaction.

ACKNOWLEDGMENT

The authors would like to thank project partners at Huntleigh Plc. and Plextec Plc., particularly S. Cook and A. Scholan, the Brunel Institute of Biotechnology, and to BT for their support and development work. Special thanks go to Prof. H. Wolff for developing the Millennium Home concept and motivating the entire team. The anonymous reviews of the paper also deserve recognition, and the authors offer their thanks to their extremely useful and pertinent comments and advice.

REFERENCES

- [1] B. Brumitt, B. Meyers, J. Krumm, A. Kern, and S. Shafer, “EasyLiving: Technologies for intelligent environments,” presented at the Conf. Handheld and Ubiquitous Computing, Bristol, UK, Sept. 25–27, 2000.
- [2] S. S. Intille, Designing a home of the future, in *IEEE Pervasive Computing*, vol. 1, pp. 76–82, Apr./June 2002.
- [3] R. Harper, *Inside the Smart Home*. New York: Springer-Verlag, 2003.
- [4] E. D. Mynatt, I. Essa, and W. Rogers, “Increasing the opportunities for aging in place,” in *Proc. ACM Conf. Universal Usability (CUU)*, Arlington, VA, 2000, pp. 65–71.
- [5] V. Stanford, “Using pervasive computing to deliver elder care,” *IEEE Pervasive Computing*, vol. 1, pp. 10–13, Jan./Mar. 2002.
- [6] S. S. Intille, K. Larson, and C. Kukla, “Just-in-time context-sensitive questioning for preventative health care,” in *Proc. AAAI 2002 Workshop Automation as Caregiver: Role of Intelligent Technology in Elder Care*, Edmonton, AB, Canada, July 2002.
- [7] A. Tinker, *Older People in Modern Society*, 4th ed. London: Longman, 1997.
- [8] “Census,” U.K. National Office of Statistics, HM Stationery Office, 2001.
- [9] A. Dowdall and M. Perry, “The millennium home: Domestic technology to support independent-living older people,” in *1st Equator IRC Workshop on Ubiquitous Computing in Domestic Environments*, Nottingham, U.K., Sept. 2001.
- [10] T. A. Salthouse, *Theoretical Perspectives on Cognitive Aging*. Hillsdale, NJ: Erlbaum, 1991.
- [11] J. Porteus and S. Brownsell, *Exploring Technologies for Independent Living for Older People: A Report on the Anchor Trust/BT Telecare Research Project*. Oxon, U.K.: Anchor Trust, 2000.
- [12] N. M. Barnes, N. H. Edwards, D. A. D. Rose, and P. Garner, “Lifestyle monitoring—technology for supported independence,” *Computing and Control Eng. J.*, vol. 9, no. 4, pp. 169–174, 1998.
- [13] B. G. Celler, W. Earnshaw, E. D. Ilisar, L. Betbeder-Matibet, M. F. Harris, R. Clark, T. Hesketh, and N. H. Lovell, “Remote monitoring of health status of the elderly at home. A multidisciplinary project on aging at the University of New South Wales,” *Int. J. Bio-Med. Computing*, vol. 40, no. 2, pp. 147–155, 1995.
- [14] J. Tomlinson, M. J. Russell, and N. M. Brooke, “Integrating audio and visual information to provide highly robust speech recognition,” in *Proc. IEEE ICASSP*, 1996, pp. 821–824.
- [15] S. L. Oviatt, “Mutual disambiguation of recognition errors in a multimodal architecture,” in *Proc. Conf. Human Factors in Computing Systems (CHI'99)*, New York, 1999, pp. 576–583.
- [16] A. Smith, J. Dunaway, P. Demasco, and D. Peischl, “Multimodal input for computer access and augmentative communication,” in *Proc. 2nd ACM Conf. Assistive Technologies*, Vancouver, Canada, 1996, pp. 80–85.
- [17] M. A. Grasso, D. S. Ebert, and T. W. Finin, “The integrality of speech in multimodal interfaces,” *ACM Trans. Comput.-Human Interaction*, vol. 5, no. 4, pp. 303–325, 1998.
- [18] S. Oviatt, “Taming recognition errors with a multimodal interface,” *Commun. ACM*, vol. 43, no. 9, pp. 45–51, 2000.
- [19] B. Suhm, B. Myers, and A. Waibel, “Multimodal error correction for speech user interfaces,” *ACM Trans. Comput.-Human Interaction*, vol. 8, no. 1, pp. 60–98, 2001.
- [20] E. Kantorowitz and O. Sudarsky, “The adaptable user interface,” *Commun. ACM*, vol. 32, no. 11, pp. 1352–1358, 1989.

- [21] D. A. Norman, *The Psychology of Everyday Things*. New York: Basic Books, 1988.
- [22] R. J. Orr and G. D. Abowd, "The smart floor: A mechanism for natural user identification and tracking," in *Proc. 2000 Conf. Human Factors in Computing Systems (CHI 2000)*, The Hague, The Netherlands, Apr. 1–6, 2000.
- [23] V. M. E. Bellotti, M. J. Back, W. K. Edwards, R. E. Grinter, C. V. Lopes, and A. Henderson, "Making sense of sensing systems: five questions for designers and researchers," in *Proc. 2002 Conf. Human Factors in Computing Systems (CHI 2002)*, Minneapolis, MN, 2002, pp. 415–422.
- [24] G. D. Abowd, E. D. Mynatt, and T. Rodden, "The human experience," *IEEE Pervasive Computing*, vol. 1, pp. 48–57, Jan./Mar. 2002.
- [25] W. W. Gaver, "Designing for homo ludens," *I3 Magazine*, vol. 12, pp. 2–6, June 2002.



Mark Perry received the B.A. degree in psychology and the M.Sc. degree in cognitive science from the University of Cardiff, U.K., and the Ph.D. degree from Brunel University, Uxbridge, U.K.

He is a Senior Lecturer at Brunel University, Uxbridge, U.K., lecturing on interactive systems design and has been involved in research into human–computer interaction since 1993. Mark has recently led several research grants in the user-centered design of ubiquitous and mobile computing and has worked as a consultant with a number of leading companies in these areas. His past research includes work on mobile medical device design, mobile collaboration, and networked appliance design, to name a few. Prior to his current position, he has been a Visiting Scholar at Stanford University, CA, and a Research Fellow at Brunel University.

Alan Dowdall received the B.Sc. degree in computer science from Brunel University, Uxbridge, U.K.

He was a Researcher on the Millennium Homes project led by Brunel University, Uxbridge, U.K. His research interests fall within human–computer interaction and inclusive design practices. His areas of expertise include dialogue design, usability inspections, and prototype development.



Lorna Lines received the B.Sc. degree in psychology from the University of Huddersfield, and the Ph.D. degree in the design of speech alarms for older adults from Brunel University, Uxbridge, U.K.

She is a Lecturer at Brunel University, Uxbridge, U.K. She has been conducting research into cognitive ergonomics and human factors of interactive systems since 1998. Prior to her current position, she has been a Researcher at Brunel University and Loughborough University.



Kate Hone received the B.A. degree in experimental psychology from the University of Oxford, U.K., the Masters degree in work design and ergonomics from the University of Birmingham, U.K., and the Ph.D. degree in the design of dialogues for speech systems from the University of Birmingham.

She is a Senior Lecturer at Brunel University, Uxbridge, U.K. She has been conducting research into the human factors of interactive speech systems since 1992 and currently leads a project looking at the human factors issues surrounding the use of emotion recognition technology.