Specification of a Decision Support System for Smart Home Applications

A case study for the graduate course
'Intelligent Decisions – Intelligent Support?'

Dennis Maciuszek

May 19, 2003

1 Decisions

1.1 Application: Smart Homes

Smart home environments are now becoming a reality. Today's pervasive computing makes it possible to seamlessly embed technology in residential facilities and automate daily living tasks the resident considers laborious.

Elderly users While young and capable people may have reasons to be skeptic towards more and more technology entering their homes, two groups that could clearly benefit from smart home applications are disabled people and, as the focus of this study, frail elderly people. Smart homes may prolong the time which a frail elderly person can continue living in his or her own home, either independently or with only light caregiving assistance.

Safety Building a smart home environment for an elderly resident usually means equipping her or his living area with technology that enables the remote and automated control of hardware, as well as monitoring resident and environment in order to detect safety risks. This study concentrates on the latter, smart homes for increasing safety.

Examples The Swedish Handicap Institute's 'SmartBo' initiative (Elger and Furugren, 1998, now called the 'SmartLab') as well as the 'Gloucester Smart House' (Adlam, 2001) realise a number of smart home functions for increasing the safety of frail elderly people. Both projects have in common that they – among other functions – have investigated

- the detection of and response to
  - unwanted water flow from taps in the bathroom,
  - unwanted heat and gas development on the cooker in the kitchen,
- as well as the automated lighting of rooms the user is about to enter.

Another typical application is the automated monitoring and operation of doors and windows; see SmartBo and the 'House_n' project (Intille, 2002).

Warnings Smart home applications supporting safety can usually detect a couple of anticipated dangers, and warn the user of these. SmartBo, for instance, shows a resident leaving the house or going to bed, if water taps and the cooker have been turned off. This is indicated by green and red lights, or for visually impaired people by corresponding tactile signals.

A service like this is certainly a good help in preventing a frail person from inadvertently causing a flood or a fire at home. Yet, it is far from being a safety guarantee. What if a physically impaired user needs to engage in a major effort, when he or she must get up from bed again to turn off a tap? What if a cognitively impaired user does not recall what the lights mean, or how to turn off the gas? Users who – consistently or occasionally – have difficulties handling the smart home's warnings may become insecure, and do the wrong thing or nothing at all.

Automation The problem of users not responding accurately to warnings may be solved by further automation. Usually, smart home systems that support safety can respond to dangers themselves. The Gloucester Smart House can automatically turn off the water tap of the bath, or the cooker in the kitchen. Both the Gloucester house and SmartBo can turn on lights. The Gloucester 'night light' for instance guides a demented user from bed to bathroom by turning on lights along the path in the appropriate sequence.

Yet, too much automation can be as bad as not enough automation. A traditional situation-response home automation system – at least one not thoroughly configured for the specific use case – bears the risk of inaccurate situation assessment or inaccurate response generation (Miller et al., 2002). For example, the automation may open a window because of high temperature in the room, neglecting the fact that there is a loud noise outside, or that somebody is smoking there, or that pollen the user is allergic to is in the air, or that it is raining. 'The sophistication of commonsense reasoning and context awareness that is required is daunting, given the current state of our understanding of these fields.' (Intille, 2002, page 81) Another aspect is that automation causes the human user to lose control, and it can discourage learning (Intille, 2002) – or encourage forgetting – in daily life.

Need for decision support How can the user remain in control, and at the same time not be left alone? Apparently, there is a need for mediation between warnings, automatic response, and the user. That is mediation between uncertainty, automated decision making, and the user. My suggestion is to provide this mediation in the form of a decision support system.

A simple example of integrated decision support would be the reminder given by the Gloucester 'bath monitor' which informs that water is running when a sensor in the bath feels water. If the tap has not been turned off after some time, the automation takes over, turning off the water (Adlam, 2001, page 23 shows architecture and algorithm in detail).

In order to solve the dilemma of the wrongly opened window, Intille suggests a light on the window that illuminates as if to say: 'It might be a good idea to open this window right now' (Intille, 2002, page 81). Based on this support, the user could then decide whether to open the window or not.

In a SmartBo environment, the user would further be able to decide to either physically walk to the window and open it, or to do so with the effort of a mouse click, thus determining the degree of automation her- or himself.

Still, such solutions appear to be scarce and somewhat limited. The 'bath monitor' reminder basically delays the automatic response, allowing the human to intervene and turn off the tap before the system. Intille's light at the window is merely an improved warning. These approaches may be sufficient for some users and for some range of situations, yet a more flexible decision support might serve more cases of usage. Further investigation into decision support in smart homes should therefore prove to be worthwile.

Scenario This text examines the issue of decision support in smart homes via a case study. A decision support system for a simple smart home scenario is to be designed, with the following technology considered given (Figure 1).

- In the bathroom, water flow from a tap and the water level of its basin are monitored. The tap can be turned off automatically.

- In the living room, the status of a window (open/closed) and the room temperature are monitored. The window can be closed automatically.

- The user is warned by the smart home software, in case the water level gets high, or the temperature gets low.

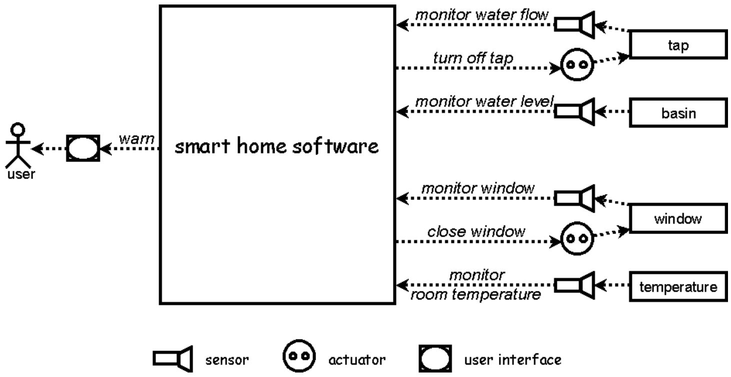

Figure 1: Given technology

The goal of introducing this technology is to increase safety of the resident by decreasing the risk of a flood in the bathroom, and of bad consequences of low temperature in the living room, e.g. getting ill.

How this technology is operated by the 'smart home software' is to be specified. Chapter 2 will fill the empty box in Figure 1 with an architecture to support decisions of the kind: How do I prevent being hurt in everyday life situations? More specifically, how should I behave, so that no flood is caused, and I do not suffer bad consequences of low temperature. At the same time, conflicting needs, such as wanting a wash-basin full to the brim, to get rid of an unpleasent smell in the living room, to sleep while the system is working, or to not be deprived of control must be considered.

1.2 Decision Making in the Smart Home

Three phases of decision making can be identified: The smart home user perceives a warning of a risky situation, processes and decides how to handle it, and then responds. These phases fit the established step-ladder model *(Rasmussen, 1979), an inductive model of human decision making (Hollnagel, 1984) comprising eight steps that form a decision making process: activation, observation, identification, evaluation, interpretation, definition of the task, formulation of the procedure, and execution.*

1. Up the ladder: Perception of a warning In order to maintain safety, the user has to notice potentially risky situations, i.e. a too high water level or a too cold room. To some extent, this is already supported by the technology given: Sensors for monitoring bathroom and living room have been installed and *activated*. They are continuously *observing* tap, basin, window, and temperature. It is then up to the smart home software to forward this information to the user, either as an assurance that the readings are okay, or as a warning that a risk has been detected. Assisted by the technology, the user *identifies* a situation requiring a decision.

2. On top of the ladder: Processing the warning After a smart home warning has been perceived by the user, a decision needs to be made about *what* to do in response. Possible alternatives are found, such as to

- ignore the safety risk, e.g. in a cold living room leave the window open because of an unpleasant smell,
- further investigate the safety risk, e.g. go to the bathroom and have a look at tap and basin,
- remove the safety risk, e.g. turn off the tap, or close the window,
- have the safety risk removed automatically by the smart home technology, e.g. have the tap turned off via the corresponding actuator,
- call for help, e.g. phone the neighbour, the caregiver, or a technician to take care of the problem,
- or to provide whatever alternative solution, e.g. open the basin drain, leave the cold room, etc.

One of the alternatives is chosen after an *evaluation* according to different criteria, and an *interpretation* regarding anticipated consequences.

Appropriate criteria here are, in order of priority, *capabilities* to actually realise the alternative, guaranteed *safety*, and finally general user *happiness*. This last criterion shall cover Intille's 'commonsense reasoning and context awareness', e.g. wanting to get rid of an unpleasant smell which the sensors cannot detect, as well as the desire to learn and be in control. General happiness is difficult to assess, yet without doubt it is also less important than safety. No alternative whose realisation would cause damage shall be chosen, not even one that causes slight damage together with great general happiness. Ignoring some water that is leaping out of the basin while enjoying a really interesting TV programme is hardly an 'intelligent decision' – at least not, if the room is vulnerable to water, or if the resident easily falls. This is consistent with Miller et al.'s 'first, cause no harm' design principle, which states that 'quality of life enhancements are desirable, but they are secondary to safety considerations' (Miller et al., 2002). Happiness should only decide between alternatives that have already been checked for the capabilities to actually realise them and the absence of threats to the resident's safety.

This view of decision making is close to Tversky's *elimination by aspects* theory (Tversky, 1972). He regards decision alternatives as sets of criteria they fulfil. The decision maker considers one criterion after another in the order of their importance, here (1) capabilities, (2) safety, (3) happiness. For each criterion, the decision maker eliminates all alternatives that do not fulfil it. The algorithm stops when only one or just a few alternatives are left. This theory explains well how one filters alternatives so that only those remain which are feasible in terms of capabilities, and which are safe. The theory is less appropriate for assessing the fuzzy criterion of general happiness, as the question if an alternative does not evoke happiness is difficult to answer.

Anticipated consequences of realising an alternative might be determined by mental simulation of the situation.

3. Down the ladder: Response to the warning After an alternative has been chosen, a decision needs to be made about *how* to put that alternative into practice so that the desired target state is reached. This involves first a *definition of the task*, e.g. closing the window, then a *formulation of the procedure*, e.g. going to the living room, climbing onto a stool, pushing the window, locking the handle, and finally the *execution* of that procedure.

1.3 Scenario Characteristics

The decisions addressed here may seem trivial with an average user in mind who is aware of the little risks of everyday life and if something happens can usually react in time. This changes when considering a frail elderly resident.

Scope Users of the scenario smart home are frail elderly people in need of extra assistance when there is no caregiver around. By definition of the situation, the user is therefore a single decision maker. This changes, when in case of a warning the user calls somebody for help. Usually, performance should improve then, depending on the competence of the assistant and the time it takes until he or she can assist – e.g. until he or she arrives at the user's home. As the user may be alone for several hours every day, it is likely that the decisions in this application happen quite regularly.

The decisions described deal with small-scale everyday life situations and problems on an operational level, i.e. no larger-scale tactical planning is required. Decisions are limited to one place – the bathroom or the living room – and involve one or only a few devices. These situation characteristics should facilitate decision making, at least if the user fully concentrates on the urgent task. Yet, many things can happen in a living environment – e.g. the doorbell might ring – which would distract or interrupt the decision maker.

Handling a safety warning often involves several incremental decisions. This applies mostly to the third phase of decision making (cf. Section 1.2). For instance, the definition of a 'call for help' task might involve choosing a suitable helper. The subsequent formulation of a procedure requires the choice of the phone number to call (home, mobile, work, . . .), then what to say, etc. This third phase can be described as problem solving, i.e. deciding what actions to take that will transfer the present state into the target state.

Time For a frail elderly person, decisions about handling safety warnings can have serious consequences. If a flood in the bathroom is not avoided, the person might slip the next time he or she goes to the toilet. Falling in old age often leads to bad injuries, further impairment of health, or even death.

The risk of an accident, e.g. a flood, increases with time. Processes like a rising water level, a dropping temperature, and a cooker heating up take several minutes until the situation becomes risky. The time period during which a risky situation can be sorted out reaches from the first perception of the warning until the last possibility to prevent an accident, e.g. until water leaps out of the basin. An accident may still be reversible though, e.g. a flooded floor can be wiped before somebody falls.

The average user can typically prevent an accident rather easily, as he or she decides and acts within a shorter time than the process takes, i.e. within seconds or only a few minutes. For those elderly users however who experience physical or cognitive limitations to their response time, the prevention of an accident is no longer trivial. When response time approaches process time, decisions may come too late.

Users The users of the smart home and its decision support system shall be frail elderly people in need of some care, but not demented or seriously ill. Safety at home does concern many elderly people. As one tends to get physically weaker in old age, the risk of certain accidents increases, and a good recovery from an injury is no longer guaranteed. Cognitive weaknesses may not be as distinct in non-demented older adults, yet if one does begin to show a slight deficit in memory, attention, or reaction time, risks will already increase. Note that statements concerning age-related limitations are always about some elderly people. That user group is in fact very heterogenous, and people differ in capabilities, interests, the environment they live in, their usage of technology, etc. This means that the need for technical aids and certain automation always depends on the individual user.

Uncertainty Physical and possibly cognitive limitations acquired in old age bring uncertainty into activities of daily living that used to be routine. Not only a demented person may some day forget water flowing from a tap, but also a person hard of hearing who no longer perceives auditory cues from the bathroom. Another example would be a person with a walking disability who does not care to get up and check the window when freezing. Automatic warnings from the smart home can reduce this type of uncertainty. Interpreting the warnings and deciding what to do about them may however bring further uncertainty. For instance, a resident new to the technology may have forgotten how to understand displayed readings, or a user who had always had her or his late spouse do the housework may panic at a warning. The job of a decision support system would be to reduce these uncertainties.

Performance Ideally, the smart home would enable the older person to live as independently as in his or her younger days. Of course, that is an abstract and perhaps utopian conception of good user and machine performance or 'intelligent decisions'. A more detailed specification of criteria is required. In the small excerpt of smart home technology chosen for the scenario, the system can be said to support its user well if no floods are caused in the bathroom and the resident does not freeze in the living room. At the same time, the user should not experience unnecessary restrictions on freedom of choice or 'general happiness'. The criteria of no flood and no freezing can be checked easily, while general happiness is subjective and difficult to measure.

2 Decision Support

2.1 Functional Requirements

The decision support system that will fill the blank box 'smart home software' in Figure 1 shall support all three phases of human decision making in the smart home context. This section lists desired support functions for each phase, supplemented with the information they will work on and the degree of automation, i.e. the workload distribution between user and computer.

1. Up the ladder: Supporting perception The decision support system monitors the sensors of the basin and the living room temperature. As long as the readings do not reach certain *thresholds*, the system assures the user that the monitored environment is safe. Once a threshold is reached, a warning is issued, together with an explanation of the situation. If possible, a potential cause of the risk – as indicated by readings from the tap and window sensors – is determined as well.

Example part 1 *Let the lowest threshold for comfortable room temperature be just above 19 °C. Normally, the system displays the temperature, and states that it is okay. Once the sensor measuring temperature indicates that it has dropped to 19 °C, a warning is issued which tells the user that the temperature is low. In addition, the sensor at the window is checked, and if it indicates that the window is open, this is given as a possible cause.*

Information this module works on are the sensor readings and the thresholds. The thresholds are being set by the user, a caregiver, or a technician during installation of the smart home or a later update. Furthermore, information about the reactive interplay of actors in the 'sensor world' home environment is needed for determining potential causes.

According to Sheridan's scale of degrees of automation (Sheridan, 2002, Table 3.2), the support specified above reaches automation degree 6 out of 8: 'The computer selects one way to do the task, executes it automatically, then necessarily informs the human.'

2. On top of the ladder: Supporting processing Whenever the smart home software has issued a warning, support of the second decision phase begins. The system internally generates alternatives regarding the handling of the risky situation. These are of the types listed in Section 1.2, but with added support like guidance to remove a risk and automated calling for help. Then, the system determines the consequence of realising each alternative by a simulation of the 'sensor world' environment and its reactive behaviour. After that, the strategy of elimination by aspects is applied. Each alternative is checked for its feasibility taking into account user and automation capabilities. It is validated according to *safety constraints*. Now, it is up to the user to judge which alternative brings the greatest 'happiness'. All alternatives that survived the elimination process are presented, ranked after previous experience with them. One additional option is always that the user tries an idea of his or her own. The user is then asked to choose one alternative. In the case that the user's decision making takes too long so that sensors already indicate that safety constraint values have been reached, the system overrides the user's authority to decide, and makes the decision itself. Of course, a 'lazy' user could always wait until this, which would unnecessarily push safety to its limits. This undesirable behaviour is inhibited by emitting regular reminders that encourage the user to come to a decision.

Example part 2 *The safety constraint shall be that the living room temperature must not sink to 17 °C. Alternatives that survive the elimination by capabilities and safety are, in order of good experience from past decisions:*

1. *The user closes the window. Consequence: It will not get colder.*
2. *The window is closed automatically via the respective actuator. It will not get colder.*
3. *The user is shown how to close the window. It will not get colder.*
4. *The user ignores the warning. It will get colder.*
5. *The user has an own idea (e.g. to turn on the heating). Consequence unknown. Further monitoring will show if the situation improves.*
6. *The user investigates the situation. The decision is postponed.*
7. *The user calls a caregiver for help. The decision is delegated.*
8. *A caregiver is called automatically (e.g. via a connected phone, or a beeper carried by the caregiver). The decision is delegated.*

The user's choice determines how the risky situation is dealt with. If the user has not chosen and notified the system by the time the sensor detects a temperature of 17 °C, the automation will proactively execute the highest ranking automatic solution, here to close the window itself (alternative 2).

The smart home software must know its own capabilities, i.e. its automation functions and what they can accomplish. It needs to store a model of its user's capabilities, i.e. what actions required in the third phase the user can perform. Furthermore, a list of potential assistants to call for help is kept. It includes people's capabilities and availability. The 'sensor world' model from the first phase is needed again to simulate reactions of the home environment to different events and interventions. Safety constraints complement the warning thresholds. Evaluations of past decisions according to resulted performance must be available so that alternatives can be ranked.

Decision making in this phase uses real human-computer collaboration. The support adapts to the user's needs by considering his or her capabilities in the suggestion of alternatives, and by finally letting the user choose the degree of automation for the following phase. A low degree of automation allows this freedom of choice, yet only until time runs out, and things get dangerous. This strategy of allowed freedom with a built-in safety mechanism would rate between 2 – 'the computer suggests alternative ways to do the task' – and 5 – 'the computer allows the human a restricted time to veto before automatic execution' – on Sheridan's scale of eight degrees.

3. Down the ladder: Supporting the response The user's decision in the second phase determines what kind of support is given in the third.

- If the user decided to deal with the situation on her or his own, no further support is given.

  Example part 3 *The user closes the window, or ingnores the warning, or has an own idea, etc.*

  This rates 1 out of 8 on Sheridan's scale of degrees of automation: 'The computer offers no assistance; the human must do it all.'

- If the user requested guidance in making the response, the decision support system generates a set of instructions to follow.

  Example part 4 *The user is given the formulation of the procedure that closes the window. Words and pictures explain how to get to the window in question, then how to reach and push it, and finally how to lock the handle.*

  Once more, the model of reactions in the 'sensor world' is accessed, as well as the model of user capabilities, so that an appropriate problem solution can be generated. The generation of instructions requires a database of explanatory texts and pictures.

  This rates 3 out of 8 on the scale of automation – 'the computer selects one way to do the task' – but with an added explanation of the task.

- If the user decided to have the automation remove the safety risk, the system determines the procedure of how to do this, and executes it by means of the smart home actuators (cf. Figure 1).

  Example part 5 *The actuator for closing the window is activated.*

  The 'sensor world' model is consulted to find the appropriate response.

  The automation rating is 6: 'The computer selects one way to do the task, executes it automatically, then necessarily informs the human.'

- If the user decided to have the automation call somebody for help, the system can automatically contact the person. If permitted by the communication means, the user can then talk to the called assistant.

  Example part 6 *The system phones a caregiver.*

  This requires an internal 'phone book' as an information source. Hardware for establishing communication was not included in the given smart home technology (Figure 1), and must therefore be added.

  Again, the degree of automation is 6, at least for the automated calling. It is however assumed that 'the computer offers no assistance' with the actual conversation, which means a rating of 1 for that task.

After the third phase of support, the made decision is evaluated according to resulting performance. If the sensor readings indicate that the risk is being removed without violating safety constraints, the decision is recorded as a good one, otherwise as a bad one.

2.2 User Interface

Interaction with the smart home must be possible from every room. There could be monitors everywhere, ideally with speakers and touch-screen functionality. Displayed information must be clearly understandable, and warnings must catch the user's attention. Figure 2 shows a possible user interface.

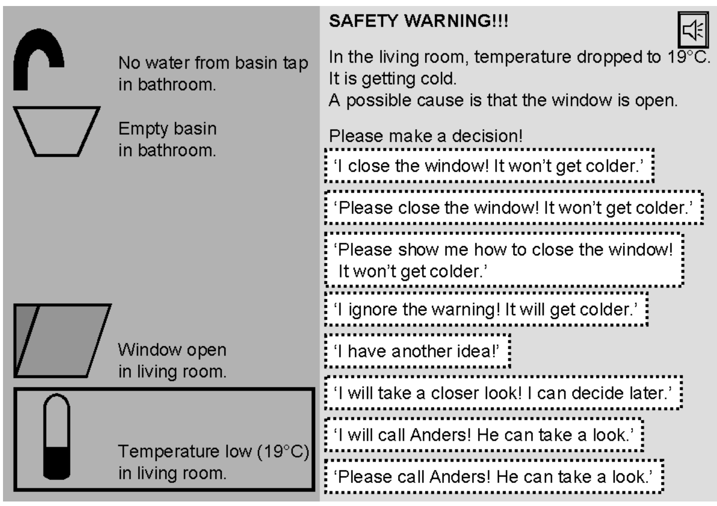

Figure 2: A user interface

Description There is only one screen, but with two areas, left and right. The left area is always visible. It displays the current status of the monitored home by icons and text. As these are constantly updated, there is no need to explicitly request safety checks. When a reading reaches a threshold and the system issues a warning, this is indicated by highlighting the related icon and text, as well as a sound signal. In addition, the right area of the screen is filled with content: The risk is named, explained, and a possible cause is given. Decision alternatives are displayed as buttons, and the user is asked to select one of them. If he or she requests further support, or readings reach the safety constraint value, a new dialogue replaces the list of alternatives. Otherwise, the right side turns blank again until the next warning.

Motivation Design of the suggested interface was guided by the intention to achieve a clear presentation of information. All sensor readings are always displayed in a non-technical fashion. Icons and natural language text serve different cognitive preferences. Moreover, these have been grouped according to the place – bathroom or living room – and cause and effect – e.g. water flow from the tap causes the basin to be filled. The list of alternatives shows what actions the system is capable of, and what consequences they effect.

Clarity of interaction is no less important. Always just one selection is to be made, and the options are clearly indicated. No advanced interaction modes put stress on inexperienced users – just clicking a button will do.

Having only one screen supports simplicity. Yet, it also makes changes less obvious. In order to cope with the phenomenon of change blindness, users with reduced vision, and anyway to catch the attention of the user, who cannot constantly observe the monitors, sound signals supplement the visual display. The highlighting in the left area, an eye-catching headline in the right, and, after all, switching from a blank half of the screen to one filled with text should further increase the likelihood that a warning is noticed.

2.3 Design

The decision support software is being designed as a multiagent system.

Architecture A multiagent system consists of a number of autonomous software modules – *agents* – that communicate and work together in order to realise the global behaviour. Multiagent systems have become an established solution to handling the complexity of a system that aims at automating processes only humans had been able to carry out before. A sophisticated smart home environment taking over many responsibilities from a caregiver would be an application of that kind. Here, an agent shall be seen as a software module which acts autonomously by reasoning on the basis of knowledge and working towards goals (deliberation), or responding to stimuli (reactivity), which communicates with other agents, and which is capable of learning.

Figure 3 presents an architecture for the smart home decision support software. It includes three types of agents: a user agent to internally represent the user, five functionality agents to realise support in the three phases of decision making – the number indicates the phase the agent is mainly active in – and several environment agents of which two represent the 'sensor world', and an indefinite number represent assisting users who may be called for help.

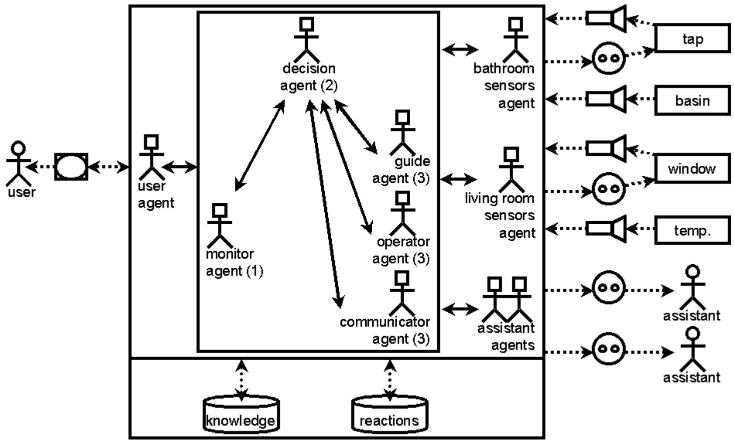

Figure 3: Multiagent architecture

The five functionality agents communicate with the user agent when they wish to collaborate with the user, i.e. by live interaction or by adaptation to user capabilities. They interact with the environment agents when they wish to read from sensors, have actuators activated – including the phone to call to an assistant – and in order to simulate behaviour of the home environment. The user agent manages the user interface, while the environment agents connect to sensor and actuator hardware. Agents access certain knowledge – e.g. the user agent knows user capabilities – and certain reactions – e.g. the bathroom sensors agent owns the reactive behaviour that the basin is being filled when water is flowing from the tap.

Agents The requirements of Section 2.1 have been divided among eight classes of agents, see Figure 4. Most classes will have only one instance, yet 2 rooms and a couple of assistants are needed. Association lines indicate which agents will communicate. Each class box contains attributes (upper part) for goals an agent is to fulfil, knowledge to base reasoning on, and reactive behaviour. Methods (lower part) will realise the functions, of which only the most basic ones are shown here. For instance, the decision agent's generation of alternatives is assumed to include the determination of consequences.

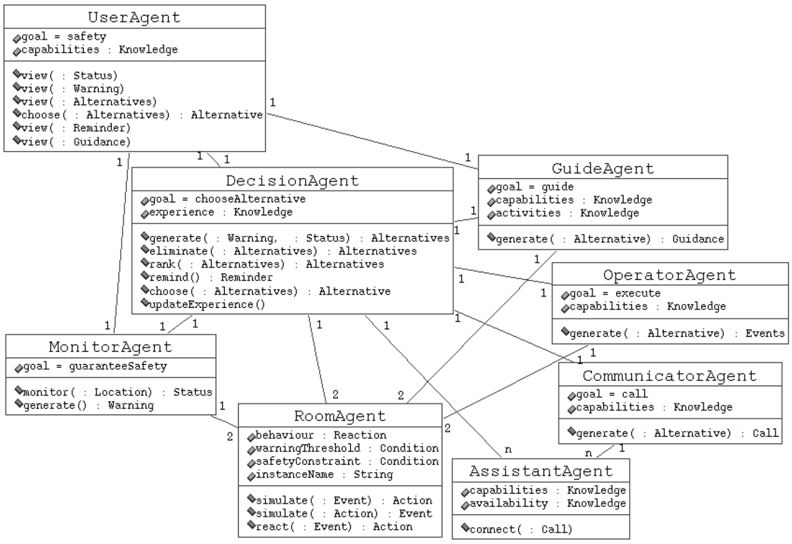

Figure 4: Agent classes

The class diagram implies a certain sequential information flow: Room agents behave, the monitor agent monitors, then the user agent and decision agent display a generated warning and support choosing an alternative, after which the guide, operator, or communicator can provide further automation. However, the agents will also continuously and concurrently be working on their own: Room agents are synchronising their behaviour with the environment, the monitor is always monitoring – forwarding warnings that the user agent displays or not, depending on current interactions. The decision agent is checking if safety constraints are still met, and several agents are simulating interventions in the environment while environment behaviour proceeds.

A much more elaborate design would be needed to really formulate a prototype system that can be implemented. Functions, agent communication, goals, knowledge, and reactive behaviour must be modelled and described in detail. This is beyond the scope of this study, whose intention was to explore the unconventional perspective of seeing a smart home increasing elderly residents' safety as a decision support system. By sketching a realisation of such support for a small scenario, it has outlined how assistance can be automated while at the same time granting the user a reasonable degree of control.

References

Adlam, T. (2001). The Gloucester Smart House project. In *HomeNet 2001*, Brussels, Belgium.

Elger, G. and Furugren, B. (1998). An ICT and computer-based demonstration home for disabled people. In *TIDE 1998 Conference*, Helsinki, Finland.

Hollnagel, E. (1984). Inductive and deductive approaches to modelling of human decision-making. *Psyke & Logos*, 5:288–301.

Intille, S. (2002). Designing a home of the future. *Pervasive Computing*, 1(2):80–86.

Miller, C. A., Haigh, K., and Dewing, W. (2002). First, cause no harm: Issues in building safe, reliable and trustworthy elder care systems. In *AAAI-02 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, pages 80–84, Edmonton, Alberta, Canada.

Rasmussen, J. (1979). On the structure of knowledge – A morphology of mental models in a man-machine system context. Report Risø-M-2192, Risø National Laboratory, Electronics Department, Roskilde, Denmark.

Sheridan, T. B. (2002). *Humans and Automation: System Design and Research Issues*. Wiley.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychology Review*, 79:281–290.