

# Smart Environments and Self-Organizing Appliance Ensembles

Thomas Heider, Thomas Kirste  
Dept. of Computer Science  
Rostock University, Germany  
{th,tk}@informatik.uni-rostock.de

## Abstract

The vision of Ambient Intelligence is based on the ubiquity of information technology, the presence of computation, communication, and sensorial capabilities in an unlimited abundance of everyday appliances and environments.

It is now a significant challenge to let ambient intelligence effortlessly emerge from the devices that surround the user in his environment. Future ambient intelligent infrastructures must be able to configure themselves from the available components in order to be effective in the real world. They require software technologies that enable ad-hoc ensembles of devices to spontaneously form a coherent group of cooperating components. This is specifically a challenge, if the individual components are heterogeneous in nature and have to engage in complex activity sequences in order to achieve a user goal. Typical examples of such ensembles are smart environments.

It will be argued that enabling an ensemble of devices to spontaneously act and cooperate coherently requires software technologies that support unsupervised spontaneous cooperation. We will illustrate why a goal based approach is reasonable and how explicit goals can be used to find system comprehensive strategies and how explicit declarative goals could be used as a benchmark to evaluate the system design.

## 1 Introduction

Driven by the vision of invisible computing, the demand for “smart” interfaces is continually growing. Such concepts are described for exam-

ple in the visions of “Ubiquitous Computing” [1] or “Ambient Intelligence” [2]. They share the belief of a smart, personal environment which characterizes a new paradigm for the interaction between a person and his everyday surroundings: Smart Environments are aware of the user and his surroundings and are equipped with computing and communication capabilities to make intelligent decisions in automated and situation-aware fashion.

However, future Smart Environments will be composed from individual components (“smart appliances”) that have to assemble themselves into a coherently acting ensemble. This requires software technologies that enable appliances to cooperate spontaneously on behalf of the user’s needs.

A rather popular scenario illustrating this application area is the smart conference room (or smart living room, for consumer-oriented projects) that automatically adapts to the activities of its current occupants. Such a room might, for instance, automatically switch the projector to the current speaker’s presentation as she approaches the lectern, and subdue the room lights - turning them up again for the discussion.

Today’s smart environments in the various research labs are usually built from devices and components whose functionality is known to the developer. So, all possible interactions between devices can be considered in advance and suitable adaptation strategies for coping with changing ensembles can be defined. When looking at the underlying software infrastructure, we see that the interaction between the different devices, the “intelligence”, has been carefully handcrafted by the software engineers, which have built this scenario. This means: significant changes of the

Table 1: *Smart Environment Projects*

Project	Intention Analysis	Strategy Generation	Strategy Source
MavHome, UTA	Learning and Prediction	Learned Procedures	Learned from User
The Adaptive House, Boulder	Learning and Prediction	Learned Procedures	Learned from User
The Aware Home, GaTech	Context Widgets; MySQL	Rule Set (manually eng.)	System Designer
Easy Living, MicroSoft	Geometry Model	Rule Set (manually eng.)	System Designer
AIRE, MIT Oxygen	Rule-based Programming	Rule Set (manually eng.)	System Designer
Intelligent Classroom, NWU	Plan Recognition	Rule Set (manually eng.)	System Designer

ensemble require a manual modification of the smart environment’s control application.

This is obviously out of the question for real world applications, where people continuously buy new devices for embellishing their home or offices. Enabling the devices to configure themselves into a coherently acting ensemble, requires software infrastructures that allow a true self-organization of ad-hoc appliance ensembles, with the ability to afford non-trivial changes to the ensemble.

Things will become even more complicated when looking at the visions of “The Invisible Computer” from Don Norman or “Disappearing Computing” from the ISTAG. This raises the following questions:

- How do you interact with smart things you are not aware of?
- How do you control devices you do not perceive?
- How to do this in a dynamic environment?

We will show that to cope with the problems of *invisible computer* and *dynamic infrastructures* we have to rely on explicit goals to allow the smart ensemble to cooperate spontaneously on behalf of the user’s needs.

The remainder of this paper is structured as follows: In Section 2 we review how present-day Smart Environment projects have addressed the above questions. In Section 3, we outline the concepts of our approach. In Section 4 we present different realizations of our approach. Section 5 will give a conclusion and an outlook.

## 2 Present-day Smart Environment Projects

A Smart Environments must identify the user’s intention and, in response, it needs to be able to generate multi-appliance strategies for a coherent ensemble reaction. It is now interesting to look at the means current projects employ for performing these obligations. Typical examples are for instance Microsoft’s EasyLiving [3] and MavHome from UTA [4].

The intelligent agents of MavHome for example predicts the inhabitant’s next action in order to automate selected repetitive tasks for the inhabitant. This prediction is based on previously-seen inhabitant interaction with various devices. In order to do this prediction the researcher of MavHome have instrumented the house with sensors and characterize inhabitant-device interaction as a Markov chain of events and utilize an *Active-LeZi* algorithm to do the prediction. To learn strategies they use a reinforcement learning agent.

The EasyLiving Geometric Model (EZLGM) provides a general geometric service for ubiquitous computing, focusing on in-home or in-office tasks in which there are input/output, perception and computing devices supporting multiple users. EZLGM provides a mechanism for both determining the devices that can be used for a user interaction and aiding in the selection of appropriate devices. The EasyLiving system has “behavior rules” that cause things to happen automatically when certain relationships are satisfied in the world model.

Table 1 summarizes the intention analysis and strategy generation mechanisms for a variety of well-known Smart Environments projects. As can be seen, there are two basic approaches to

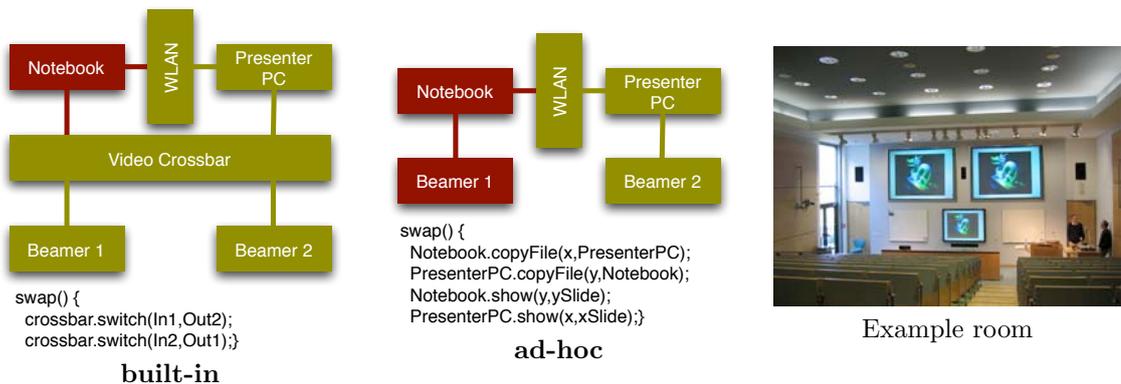


Figure 1: Achieving the same effect with different ensembles

strategy generation: (a) learn from user – by observing the user’s interaction with the infrastructure, as is done by MavHome (b) learn from system designer – by receiving a set of behavioral rules, as has been done for EasyLiving.

Unfortunately, both approaches are not viable any more, as soon as we look at dynamic ensembles. Consider the example outlined in Figure 1, which shows the built-in infrastructure of two hypothetical conference rooms is displayed (greenish boxes). The room at left provides two beamers and a video crossbar, enabling a rather straightforward way for swapping two presentations. At right, the conference room just contains a single beamer; the second one has been presumably provided by an attendee. (In both sketches, the reddish boxes denote components that have been added dynamically.)

Clearly, both conference rooms require two significantly different strategies for realizing the user’s goal of swapping two presentations. And, while in the built-in case one maybe could expect the room designer to provide a suitable macro, this is not realistic for the ad-hoc situation: No designer of a smart room can be expected to anticipate every possible ad-hoc extension of the built in infrastructure and to provide control strategies for every possible activity that could be performed with the thusly extended ensemble. Therefore, approaches such as EasyLiving are not viable for the case, where the environment’s capabilities are provided by a dynamic ensemble.

On the other hand, the approach taken by MavHome, to learn strategies from the user, is

not an option either: If a substantial set of devices is *invisible* to the user, they can obviously not become part of a control strategy the user might develop. Therefore, a system can not learn from the user how and when to use these devices.

Since in dynamic ensembles neither system designer, nor system user have an overview over the complete ensemble and its potential, there is no human being that could provide strategies to this ensemble.

Either, the user has to be made aware of the available devices and their potential (pushing the responsibility back to the user), or the ensemble *itself* must become able to develop strategies on its own, based on the user’s objectives. With respect to this, it should be noted that the systems developed in the above projects have *no* explicit notion of the user’s objectives: They learn *procedures* from the user (or receive them from the system designer), but they have no concept of the *effect* of these procedures with respect to the user’s objectives.

The consequences of the Ubicomp Visions are:

- *Disappearing Computer*: The system can’t learn the strategies from the user!
- *Computers are everywhere - dynamic ensembles*: The system designer can’t predict the ensemble and therefor can’t provide predefined system strategies!
- *So*: Strategies have to be generated dynamically by the ensemble!

### 3 Goals

The consequence of disappearing computer and dynamic ensembles is that we need appliances that cooperate spontaneously and are able to *autonomously* generate strategies that accomplish the goal of the user. To make that possible we rely on goal based interaction.

#### 3.1 Goal based Interaction

When people are using their technical infrastructure they have certain goals they want to achieve; a certain satisfaction they want to experience. This goal-driven nature of human activity is a well-accepted concept in cognitive psychology. However, today's technology-rich environments force us to think of interaction in terms of the individual "functions" that the numerous devices provide: functions such as "on", "off", "play", "record", etc.. When interacting with devices, we select, parameterize, and then execute functions these devices provide. Upon execution, they cause an effect: a broadcast is recorded on videotape, the light is turned brighter, and so on.

But then, a user is not really interested in the function he needs to execute on a device - it is rather the function's effect which is important.

This observation immediately leads to the basic idea of goal-based interaction. Rather than requiring the user to invent a sequence of actions that will produce a desired effect ("goal") based on the given devices and their capabilities, we should allow the user to specify just the goal ("I want to see 'Star Wars' now!") and have the ensemble fill in the sequence of actions leading to this goal. Goals allow services to be named by their semantics - i.e., by the effect they have on the user's environment - thereby evading the problems of syntactical service addressing.

Goal-based interaction requires two functionalities: **Intention Analysis**, translating user interactions and context information into concrete goals, and **Strategy Planning**, which maps goals to (sequences of) device operations (see Figure 2).

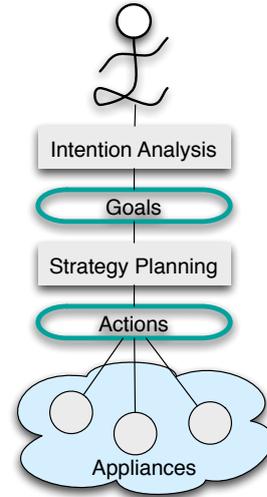


Figure 2: Principle of goal based interaction

#### 3.2 Explicit Goals

In order for a system to *autonomously* generate strategies for achieving certain goals, we need a mechanism to *explicitly* represent goals, which allows the system to reason about goals and different ways for achieving them. Specifically, we need a *declarative* representation of goals. The application area Smart Environment is concerned with achieving effects of interest to the user in his current environment - therefore, an explicit representation of user goals is basically given by a suitable state model for the environment.

Goals are then represented by state vectors that are to be made true in the given environment.

Once the concept of explicit declarative goals is made available, it serves *two* purposes: (a) it allows goal-based interaction for the user (b) it is the foundation for the autonomous computation of control strategies by the ensemble.

So, we are no longer restricted to strategies that are either learned from the user or from the system designer - the ensemble itself is leveraged to unsupervised spontaneous cooperation.

Following, we give two examples for the explicit definition of goals and the accompanying strategy computation mechanism.



Figure 3: Goal-based ensemble control: Example

## 4 Exploiting Goals

### 4.1 Achieving Effects

Once an explicit declarative representation of the user goal is available, it becomes possible to exploit partial-order planning mechanisms. This requires to describe the operations provide by the available devices as *precondition/effect rules*, where the preconditions and effects are based on the environment state model. These rules then can be used by a planning system for deriving strategies for reaching user goals, which consider the capabilities of all currently available devices. The planning system receives the goal identified by the Intention Analysis. It must then find a strategy that changes the environment from its current state to the goal state. This can be understood as a classical planning problem:

- The goal is given as a set of positive and negative literals in the propositional calculus.
- The initial state of the world (resp. the state of the system and the environment-condition which is known to the system) is also expressed as a set of literals (fact file).
- The actions provided by the available devices (“operators”) have to be characterized using a suitable definition language. It describes the action’s relation to the environment: it contains a set of preconditions that must be true before the action can be executed and a set of changes, or effects, that the action will have on the world. Both the

preconditions and effects can be positive or negative literals. The modelling of the device functions as planning operators is most important: it is the base for providing goal-based interaction with a dynamic infrastructure.

The critical aspect here is the expressive power of the model used for describing device operators, which needs to be strong enough to capture at least the operational semantics of today’s consumer appliances.

#### 4.1.1 Concrete example

As example, consider the situation outlined in Figure 3, left, where a user would like to increase the brightness of his TV set. Assuming the TV is already set to maximum brightness, the sensible reaction of the ensemble would be the one given at right: reduce ambient light. In order for an ad hoc ensemble to arrive at this conclusion, TV set, lamp, and shutter must provide a description of their capabilities, similar to the one given below<sup>1</sup>:

```

;lamp's action
(:action turn-down :parameters (?l - lamp)
 :precondition (luminosity ?l high)
 :effect (and (not (luminosity ?l high))
              (luminosity ?l low)))

;shutter's action
(:action close :parameters (?s - shutter)
 :precondition (and (time day) (open ?s))
 :effect (and (not (open ?s)) (closed ?s)))
(:axiom :vars (?s - shutter)
 :if (or (time night) (closed ?s))

```

<sup>1</sup>For sake of brevity, this capability definition has been very much simplified.

```

:then (luminosity ?s low))
;tv set action
(:action turn-brighter :parameters (?b - tv-set)
 :precondition (or (< (brightness ?t)
                    max-brightness)
                  (forall ?x (luminosity ?x low)))
 :effect (and (= (brightness ?t) max-brightness)
              (brighter achieved)))

```

Then, based on a specific situation given by

```

(:objects l - lamp s - shutter - t tv-set)
(:init (= brightness t max-brightness) (time day)
        (open s) (luminosity l high))

```

a suitable plan for the goal (`brighter achieved`) could then be computed as (`(close s) (turn-down l)`).

#### 4.1.2 Requirements for Planning Tools

The experience from the modelling of our domain has shown that we need a planning environment that supports conditional effects and disjunction in the preconditions – this allows a compact representation of device operator sets. Furthermore it is mandatory to have universal quantification in the preconditions and the effects. This for instance allows to define operators that apply to an arbitrary number of objects – which is extremely important in an environment that is dynamically extensible.

For our first running prototype we used the UCPOP planner. But the experiences has shown, that the expressiveness of this systems operator definition language was not well suited for modeling various problems of our application domain. Especially the feasibility to modeling temporal and continuous processes was missing, what is necessary to provide a reasonable time and resource management. Important is also to be able to representing mixed discrete/continuous domains. Amongst other approaches from different researchers the Planning Domain Definition Language (PDDL2.1) goes in this direction. In the current implementation of our planning component we use the systems Metric-FF [5], LPG and MIPS. Our default system is Metric-FF. Should it find no solution for the given problem, automatically one of the other systems gets the task to find a solution. See [6] for details. Our next step is to distribute the strategy finding

among all available appliances or devices. Therefore we currently evaluating different strategies for distributed planning.

## 4.2 Resource Scheduling/Optimization

Another important application area is the definition of an optimal ensemble behavior regarding the mapping of (sub)-tasks to available resources. Here, a metric has to be defined that describes, how “good” a certain mapping (“schedule”) of tasks to the available resources is and which allows to compare different schedules with respect to their optimality. So, these types of goals provide an explicit statement of a system designer’s idea of optimal ensemble behavior<sup>2</sup>. This can be regarded as a theory of optimal ensemble behavior. The ensembles responsibility w.r.t. unsupervised spontaneous cooperation is then to jointly approximate this global optimum as good as possible. These implicit goals will be triggered by the situation through the intention analysis and must then be achieved by the appliance ensemble. In this section we present an example scenario of a multi-display environment.

In meeting scenarios, where many people come together, many documents have to be visible for the participants. The accessing of digital information in environments like meeting rooms is increasingly supported by new available technologies. For example the concept of projecting various content on different surfaces in a real environment including walls, floors and desks has been proposed recently. Pinhanez for example introduced an interesting method for a ubiquitous display called Everywhere Display [7] which uses a projector with a pan-tilt mirror for covering a large area with a single or a few projectors. What the available technologies are lacking of, is a concept of what information (document) should be displayed on which display (“Display Mapping”).

In our scenario we have a meeting room with an ad-hoc ensemble of different projectors, steerable projectors, electric screens, notebooks and a number of users (see Figure 5 to get an idea).

<sup>2</sup>Interestingly, in contrast to the user goals discussed in the previous section, scheduling objectives tend to be *predefined*, implicit goals.

$$q_{\max} = \max_{\substack{ym \in YM \\ dm \in DM}} \left( \sum_{\substack{u \in User \\ d \in Document}} i(d, u) * \max_{y \in dm} v(ym \ y, u) * p(y, ym \ y) \right) / \left( \sum_{\substack{u \in User \\ d \in Document}} i(d, u) \right)$$

Figure 4: Maximum Quality Function for a Multi-User Multi-Display Environment

Now, it is interesting to provide an explicit representation of such an ensemble’s optimal behavior with respect to the Display Mapping. This representation is given by Fig. 4:  $q_{\max}$  provides a *global* and *generic* definition of an ensemble’s optimal behavior. Once  $q_{\max}$  is given, it allows *any* ensemble to optimize its Display Mapping.

This specific formulation of  $q_{\max}$  is based on the concepts of *Visibility / Projectability* and *Importance*. Obviously, this definition of  $q_{\max}$  represents the idea that all participants of the meeting can see the documents that are important for the respective participant in the best possible way.

The Visibility value results from the angle between the view direction of the user and the respective screens:

Visibility  $v : Surface \times User \rightarrow [0; 1]$ , *e.g.* simplified as:

$$v(s, u) = \max \left\{ 0, \frac{\langle \vec{n}_s, \vec{u} - \vec{s} \rangle}{\|\vec{u} - \vec{s}\|} \right\}$$

The Projectability value results from the angle between the projection direction and the respective (electric) screens:

Projectability  $p : Display \times Surface \rightarrow [0; 1]$

The Importance value depends on the agenda of the meeting and the role of the user in the different situations:

Importance  $i : Document \times User \rightarrow [0; 1]$

The goal is the optimization of the function  $q_{\max}$  in Figure 4 which results in a Display Map, which associates each display (*e.g.* a steerable projector) with a display surface (*e.g.* a projector screen) and a Document Map, which associates each document with a display:

Display Maps:  $YM = Display \rightarrow Surface$

Document Maps:  $DM = Document \rightarrow \mathbb{P}Display$

Note that the definition of  $q_{\max}$  in Fig. 4 is obviously not complete – for instance, it

lacks a notion of “history”, that would keep a document from confusingly hopping from display to display as the user slightly shifts position. However, the objective here has not been to develop the *optimal* definition for  $q_{\max}$ , but rather to argue that *explicit* global optimization goals are required in order to allow dynamic ensembles to cooperate. Only goal definitions that completely abstract from the concrete ensemble composition – such as  $q_{\max}$  – enable *arbitrary* ensembles to optimize their behavior.

### 4.3 Experimental settings

The two pictures in Fig. 5 show the behavior of a typical ensemble controlled by  $q_{\max}$ . In the left image of Fig. 5 we have a scenario with one steerable projector, one screen and two users, whereby the right user gives a presentation. After adding a second beamer and two notebooks the system automatically calculates a remapping of the document display assignment (based on the maximum quality function of Figure 4), which you can see in the right image.

We have tested this scenario in our Environment Simulation System, a visual simulation tool for Smart / Instrumented Environments that provides a simple rendering & physics simulation server, to which device and strategy planning agents may connect via sockets. The behavior of sensors, devices, simulated users, environment geometry, etc. can be controlled and changed dynamically.

Our test settings up to this point seems to justify the assumption that our intuitive formulation of a globally optimal ensemble behavior w.r.t. display mapping is heading in the right direction. Now it makes sense to look at the actual aspect: unsupervised spontaneous cooperation regarding joint approximation of this goal. Specifically, this goal-based approach now allows us to evaluate *different* distributed ad-hoc cooperation strategies w.r.t. their fidelity in ap-



Figure 5: Environment Simulation System

proximating the global optimum. Currently we testing distributed versions of the well known algorithms “local beam search”, “hill climbing” and “simulated annealing”. We consider market-based strategies for distributed cooperation as a promising approach.

## 5 Conclusion and Outlook

This paper describes how we can deal with the problems of *invisible computer* and *dynamic infrastructures* if we rely on explicit goals to allow the smart ensemble to cooperate spontaneously on behalf of the user’s needs. To make goal based interaction possible we need to have an explicit State Model of the environment. This is the foundation for the formulation of explicitly uttered goals by the user and for the formulation of implicit persistent goals.

Currently, we are working on an unified environment ontology, which can be used to formulate both described kinds of goals.

**Goal function as Benchmark:** As different software infrastructures emerge, criteria are required by which the potential and efficiency of different solutions can be compared. Clearly, user trials - the ubiquitous evaluation strategy for pervasive computing applications - are not a viable approach in this case: the users of system software are application designers. Doing extensive user trial with highly trained experts is prohibitively expensive.

Therefore, it seems desirable to identify comparison criteria that can be evaluated at a formal level, using a standardized set of example prob-

lems - a set of Benchmarks.

It may be interesting to use explicit definitions of optimal behavior as a means for creating such benchmarks. So, as already outlined above, different approaches to computing display mapping could be compared with respect to their ability to approximate  $q_{\max}$  (or a more refined “theory of an optimal display mapping”).

**To summarize:** Smart environments promise to enable ubiquitous computing technology to provide a new level of assistance and support to the user in his daily activities. An ever growing proportion of the physical infrastructure of our everyday life will consist of smart appliances. In our opinion, an effective realization of smart environments therefore inherently requires to address the challenge of unsupervised spontaneous cooperation for ad-hoc ensembles of smart appliances.

We argue that a possible solution should be based on the fundamental concept of goal based interaction, because this enables an ad-hoc ensemble to generate strategies, instead to be dependent on predefined or learned strategies.

## References

- [1] M. Weiser. The computer for the 21st century. *Scientific American*, 3(265):94–104, 1991.
- [2] E. Aarts. Ambient intelligence: A multimedia perspective. *IEEE Multimedia*, 11(1):12–19, 2004.
- [3] <http://research.microsoft.com/easyliving/>.
- [4] <http://mavhome.uta.edu/>.

- [5] J. Hoffmann. Extending FF to numerical state variables. In *Proceedings of the 15th European Conference on Artificial Intelligence (ECAI-02)*, pages 571–575, Lyon, France, July 2002.
- [6] T. Heider and T. Kirste. Supporting goal-based interaction with dynamic intelligent environments. In *Proceedings of the 15th European Conference on Artificial Intelligence*, pages 596–600, Lyon, France, July 2002.
- [7] C. Pinhanez. The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces. In *Proc. of Ubiquitous Computing 2001 (UbiComp'01)*, Atlanta, USA, September 2001.