Ambient Compass:
One Approach to Model Spatial Relations

Petr Aksenov, Geert Vanderhulst, Kris Luyten, and Karin Coninx

Hasselt University – transnationale Universiteit Limburg
Expertise Centre for Digital Media – IBBT
Wetenschapspark 2, 3590 Diepenbeek, Belgium
`{petr.aksenov, geert.vanderhulst, kris.luyten, karin.coninx}@uhasselt.be`

Abstract. The knowledge of spatial arrangements of objects is an important component for the design of migratable user interfaces that target pervasive environments**.** Objects in these environments are often moving around individually, which leads to a highly dynamic and unpredictable environment. Due to its nature, spatial information cannot be described exhaustively, and uncertainty and imprecision need to be taken into account during both the design phase and at runtime. We present an approach to model dynamic spatial information, providing it with the ability to interpret to some extent uncertain and imprecise knowledge. We then integrate this type of spatial-awareness into *ReWiRe*, a framework for designing interactive pervasive environments, in order to improve its user-interface distribution techniques.

1 Introduction

Spatial information has been named one of the most important knowledge about an environment in a large amount of research (e.g. [7,16]). Nowadays, both indoor and outdoor environments are getting populated with mobile devices, implying that many interaction resources are carried around by users most of the time. The position and orientation of these devices can have an impact on the way users interact with their surroundings; for instance, it is more likely that users will execute a task by making use of resources in their vicinity. Therefore modelling the spatial arrangements of computing devices is an important step to get insight into the full topology of the environment. However, it is still unclear how spatial information should integrate with existing models.

In this work we present an approach to model the spatial behaviour of an environment, i.e. positional relationships that apply between objects, providing it with the ability to reason about uncertain and imprecise spatial knowledge. This type of knowledge allows that spatial arrangements between interacting resources are considered in a more natural way, thus aiming to improve the overall spatial-awareness of the environment. We propose a model, denoted as *The Ambient Compass*, to capture spatial knowledge (section 3) and discuss its implementation (section 4) and integration into a pervasive computing framework (section 5). But before we briefly present and discuss some related work that exists in the field.

2 Related work

Elaborating the form of knowledge about location is often an inherent part of the engineering process in the pervasive computing area, and a great deal of research has been devoted to this problem. For example, Bandini et. al [4] present the "commonsense spatial model" based on two concepts of "place" and "conceptual spatial relation". They then use their notions to discuss about a possible reasoning technique over the model in a way that a high-level understanding of the situation can be obtained from a combination of initial factors. Another very interesting location model is described by Satoh in [17], where special emphasis is put on the fact that pervasive environments are addressed. The model is presented as a general purpose one intended to address problems of managing location-based services. The underlying principle of the model is in dealing with virtual counterparts that are digital representations of the actual physical objects and spaces. In both works the concept of an object's closeness area is introduced in either form, but no expansion on the topic is given. Kortuem et. al [14] deal with the problem of utilising spatial information in creating new types of user interfaces and use a graph to model spatial arrangements of the system. The graph represents the spatial infrastructure of a system at a certain moment of time, and the system is then described over time by a sequence of these graphs. However, the question of uncertainty is excluded from the discussion in all of the works mentioned, but is an important part of our approach.

The concept of fuzziness is very much discussed in almost every area of research where spatial information is involved. Guesgen [9] showed the possibility to introduce fuzziness into spatial relations in general, though with a conclusion that the actual implementation of fuzziness depends on the model chosen for representation. A thorough overview of ontological modelling of spatial information as well as an extensive discussion of possibilities to implement fuzziness in such an ontology is given by Hudelot et al [13]. The work also contains an excellent collection of references to other publications in the field. Despite the fact that the application area of the presented ontology is within the image interpretation, the ideas presented can serve as a good source of information for extending both the number of concepts in our ontology and the fuzzy elements therein. Apart of the spatial information, several example of incorporating fuzziness into ontologies, from a simple and direct implementation [8] to a detailed and theoretically supported analysis of the problem as a whole [5,18,19], have been published quite recently, thus indicating that this is still a topical problem with promising trends.

3 Modelling the Compass

The model defines basic concepts natural in spatial structures. This includes positioning information, orientation angles and a division of space into the "hasOnLeft", "hasOnRight", "hasBehind", and "hasInFront" relationships that can take place between two interacting resources. The proposed classification aims at giving an application the possibility to speak a language similar to that of humans when they talk about spatial arrangements. In general, there should be two more relations, "hasBelow" and "hasAbove", but due to the increasing complexity their consideration has been postponed. Various techniques to obtain information about the location and orientation of an object from sensors exist [11,12,15]. Using this information, we can update our model in real-time and derive spatial relations by dividing the space around a resource into eight zones as depicted in Fig.1. Either one relation (e.g., "hasOnRight") or two relations (e.g., "hasOnRight" as well as "hasInFront", or, similarly, "hasOnRight" as well as "hasBehind") can apply between two resources at a given moment.

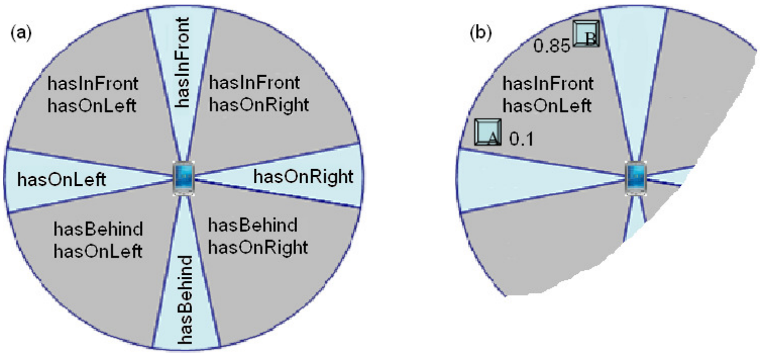

Fig.1. (a) *The Ambient Compass* divides the space around a resource into eight zones; (b) resources belonging to the same zone of the compass are distinguished by means of assigning each of them a degree of membership to this zone.

This division may remind of the way we generally refer to the parts of the world: north, east, south, west, north-west, north-east, south-east, and south-west. The boundaries of each of the main four zones are two rays symmetrically drawn to the left and right, or to the above and below, respectively, of the corresponding axis line. The slope angles of the rays depend on the resource in question therefore a set of experiments will be required to decide upon the best strategy to set them for different groups of devices. Additionally, each relation of type "has" has an inverse relation of type "is", so that if, for example, device *D1* "hasOnLeft" device *D2* then device *D2* "isOnLeftOf" device *D1*. This makes our model smoother and also simplifies queries executed on the ontology.

This simple model acts as a basis on which we build two extensions to get a more extensive model capable of handling relevant uncertain and imprecise knowledge about the spatial world.

3.1 Adding fuzziness

The first extension deals with ambiguity which appears when relations are determined. In Fig. 1b), devices *A* and *B* both belong to the "hasInFront-hasOnLeft" area of the central device, but it is obvious that their actual position with respect to this device is different. Therefore treating *A* and *B* as spatially equal would be erroneous. A possible solution to this issue lies in introducing the concept of fuzziness into establishing the four relationships. Several ways of extending ontologies with fuzzy information exist [5,8,18,19] and all of them deal with introducing in either way a degree of membership of each individual and/or relation to a certain domain. In the case of the ambient compass, these domains are its eight zones, and in Fig. 1b) you can see devices *A* and *B* have different values (0.1 and 0.85, respectively) of being in front of the central device. The membership of the second relation in the zone is such that the sum of the two values equals 1. Incorporating this extension allows to keep the relationships as appropriate using the corresponding weights, thus providing desired flexibility as well as a sort of precision in defining a more truthful type of the actual relationship between two resources. The same idea applies to the zones where there is only one relation, with the weights standing for closeness of the object in question to the corresponding adjacent zone with two relations.

3.2 Defining nearby regions

The second extension results primarily from the way how humans perceive spatial information. The concept of closeness of one object to another usually varies depending on the number of factors one considers to matter in a given situation and has already been pointed out as a subject of special attention in a number of research in either form [4,17]. The solution we suggest consists of two parts. The first one defines the concept of the "nearby" spatial relation for different types of interaction resources present in a pervasive environment. The second one defines reasonable spatial regions for each different type of interaction resource – with the resource itself being the central point – within which the corresponding "nearby" relation can be established between the resource in question and other resources. We plan to involve the concept of fuzziness into the definition of "nearby", too. This means that there is no strict division into "nearby" and "not nearby", but a degree of how much an object is "nearby" is used instead and is represented by a real number in the range [0;1]. Considering distances in such a way can provide solutions in situations where no perfect match can be found but still a positive response can be obtained. Some good examples of this approach are given, for instance, by Guesgen [10].

In addition to identifying the nearby areas, we also try to predict the behaviour of interaction resources by means of analysing their previous behaviour and reasoning over the current corresponding spatial relations between them. We use the concepts of *device availability function* and *device importance* introduced in [3] to address it.

4 Engineering the compass

Information about the model is presented in the form of an ontology and is therefore a set of concepts and properties that relate these concepts to each other. The ontology is created in the OWL language using the Protege-OWL editing tool [1]. The choice of OWL was directed by the latest trends in the development of the semantic web world

[2], following the endorsement of the W3C organisation[1]. A part of the ontology is shown in Fig.2. Due to the highly-dynamic nature of targeted environments, we do not instantiate any resources during the design phase – all instances are created at run-time.

To deal with fuzziness, we decided to use the approach suggested by Gu et. al [8] due to its self-evidence and computational simplicity. We extended our ontology with class "FuzzySpatialRelation" that has two object properties, according to the number of interacting resources involved, and two data-type properties that keep the fuzziness values (see Fig.2). It is important to note that two values are used to define fuzziness since the relations between two resources are possible in both directions. This allows having only one instance of the "FuzzySpatialRelation" class for two compasses (one from each interacting resource), keeping the entire ontology simpler.

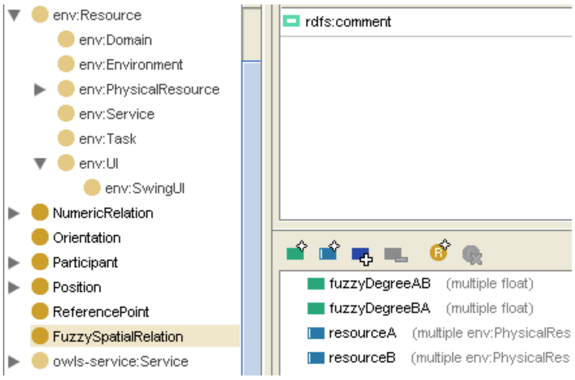

Fig. 2. Ontology as it appears in Protege-OWL. The highlighted "FuzzySpatialRelation" class on the left extends the basic concepts of the spatial model with a possibility to consider uncertain knowledge by means of introducing two additional numerical properties.

Localisation systems have a certain accuracy of measurements, resulting in a difference between the predicted and the actual location [6]. Paying our attention to this important factor, we consider a possible actual location of a resource in such a way that if the measured location has produced the fuzziness value of 0.8 then the relation's membership could vary from (0.8-δ) to (0.8+δ), where δ is the allowance parameter. The value of this parameter depends on the quality of the measuring equipment and must be determined empirically. This structure can also be used in the case of the zones with one relation. Here, the data-type property would stand for the closeness of the resource to one of the zone's borderlines so that two objects standing on the opposite sides of the zone would have different fuzziness values in the corresponding "FuzzySpatialRelation" component of the spatial ontology. The same can apply to the areas surrounding the borderlines of the zones. In other words, when it is not clear whether the position of the resource better corresponds to being on just one side of the other resource (e.g., only "hasOnRight") – as per the calculations – or should the second relation be hypothetically considered as well (e.g., both "hasOnRight" and "hasInFront") because of the measurements.

---

[1] http://www.w3.org

5 Using the Compass

5.1 Integration with *ReWiRe*

In order to validate the proposed spatial model, we have integrated it in *ReWiRe*, a framework to design interactive pervasive computing environments [20]. In *ReWiRe*, an environment is described using an upper ontology that includes concepts to represent generic resources found in a pervasive computing environment such as users, computing devices, services, tasks, etc. Aggregated with the framework's upper environment ontology, our model provides the spatial context of resources in the environment, which is for instance exploited to improve a distribution algorithm for user interfaces amongst multiple screens as well as to improve location-awareness of users in (unfamiliar) computer-augmented environments. Another important intention behind the integration of the proposed approach in the framework is to allow the use of the same language during both design and run-time phases. This assures that all changes happening during run-time can be interpreted by the designer in the way they are used at the design step.

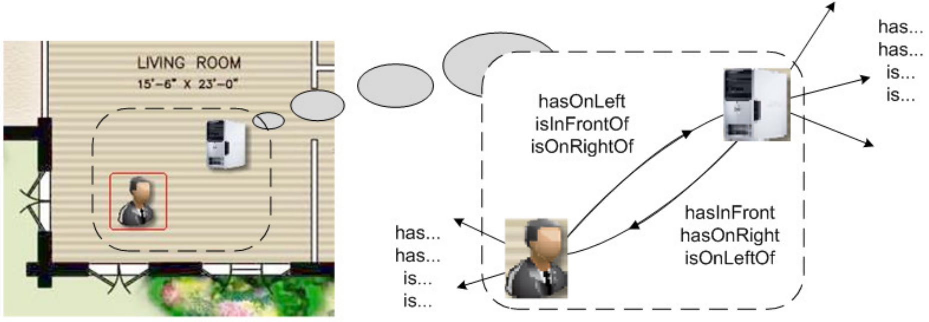

Fig. 3. Visualisation of the environment in *ReWiRe*. Resources have a set of spatial relations between them which always exist in "has+is" pairs. When a change to the location or orientation of any object happens, new spatial relations for this object are derived.

Fig. 3 shows a plug-in for *ReWiRe* we designed that simulates the movement and rotation of physical objects. Objects are overlaid on a map that represents the environment and any action executed on this view (e.g. dragging a resource) results in an update of the underlying model, in particular of its spatial relations. Besides, updates in the model triggered by sensor readings can be observed using the tool along with their spatial impact on other resources.

5.2 Action Scenario

As a brief illustration of possible use of the approach, let's consider a room with two vertically positioned displays.

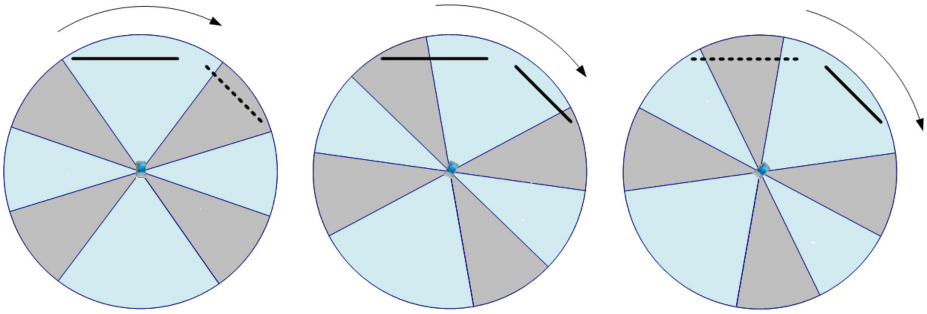

Fig. 4. The displays change their status from inactive (dashed line).into active (solid line), and vice versa, in response to the PDA turning clockwise. The change of the active display is preceded by the state when the image is shown on both of them.

Assume that a user with a PDA is oriented towards the left display and is projecting something from the PDA on it. The right display is inactive which is indicated by a dashed line (Fig. 4). At a certain moment, the user begins to turn clockwise so that the right display moves – from the PDA's perspective – from being equally "isInFrontOf" and "isOnRightOf" to much more "isInFrontOf"; whereas the left display starts holding both "isOnLeftOf" and "isInFrontOf" relations. When the PDA's rotation reaches a certain angle, the image of the PDA is copied to the right display, activating it (the dashed line becomes solid), but still being shown on the left one as well. Having observed that the PDA keeps turning, the compass discovers that the left display, though still staying in front, is already considerably to the left of the PDA and therefore can be released (the solid line becomes dashed). This example, in particular, illustrates how this kind of spatial awareness can also be used to smooth the procedure of redistributing a user interface.

6 Discussion

We presented an approach to describe spatial information intended to address pervasive environments. Its main advantages are 1) relative simplicity, a crucial factor in dealing with highly dynamic pervasive environments; 2) human friendliness, i.e. an easily recognisable interpretation of this kind of information by humans who have become a considerable part of pervasive environments; and 3) an ability to handle uncertain, incomplete knowledge which is also natural to pervasive structures. Due to being part of the *ReWiRe* framework, our approach is meant to assist designers of user interfaces in the domain of pervasive environments in general rather than in a specific type of applications.

Implementation of the underlying structure of the compass as part of *ReWiRe* has been completed and its visualisation is currently in progress. In addition, the short-term future development and improvement of *The Ambient Compass* includes, first of all, the elaboration of the concept of "nearby" for different groups of interacting resources. In particular, it will take into account their geometrical sizes. The current division into eight zones, as well as considering only four different relations, is a straightforward viewpoint. However, a modified (e.g. asymmetrical) version of the division might suit better for certain tasks, or a more precise subdividing – into more zones – might be necessary. Since no actual evaluation of the current version has been completed, discussing about these further possibilities becomes somewhat unfounded. Therefore one of the early things we plan to do is validate the current version of the approach in an experimental setup on a set of appropriate user interface distribution tasks. Based on the results, we will get a more sophisticated view of the model and will have clues on its amendment and means to improve the algorithms. In particular, on correcting values of allowance parameter δ and rays' slope angles for different devices and in different situations. For example, it is very likely for the slope angle to be a function of distance between the central object and the targeted resource. Another possible useful extension that comes directly from the above discussion about the layout of the zones is to give the designer an ability to define the zones manually, as appropriate for a task. In the long run, we consider extending the compass to the third dimension by means of introducing the "hasBelow" and "hasAbove" spatial relations.

Acknowledgments. Part of the research at EDM is funded by EFRO (European Fund for Regional Development) and the Flemish Government. Funding for this research was also provided by the Research Foundation -- Flanders (F.W.O. Vlaanderen, project CoLaSUE, number G.0439.08N).

References

1. Protégé-OWL, http://protege.stanford.edu/overview/protege-owl.html
2. Web Ontology Language, http://www.w3.org/2004/OWL/
3. Petr Aksenov, Kris Luyten, and Karin Coninx. Reasoning Over Spatial Relations for Context-Aware Distributed User Interfaces. In Anders Kofod-Petersen, Jörg Cassens, David Leake, and Marielba Zacarias, editors, Proc. of 5th Int. Workshop on Modelling and Reasoning in Context (MRC 2008), pp. 37–50. TELECOM Bretagne, June 2008.
4. Stefania Bandini, Alessandro Mosca, and Matteo Palmonari. Commonsense Spatial Reasoning for Context-Aware Pervasive Systems. In Thomas Strang and Claudia Linnhoff-Popien, (eds.), LoCA 2005, LNCS, vol. 3479, pages 180–188. Springer, 2005.
5. Silvia Calegari and Davide Ciucci. Integrating Fuzzy Logic in Ontologies. In Yannis Manolopoulos, Joaquim Filipe, Panos Constantopoulos, and José Cordeiro, (eds.), ICEIS 2006, pp. 66–73, 2006.
6. David Dearman, Alex Varshavsky, Eyal de Lara, and Khai N. Truong. An Exploration of Location Error Estimation. In John Krumm, Gregory D. Abowd, Aruna Seneviratne and Thomas Strang (eds.), UbiComp 2007, LNCS, vol. 4117, pp. 181–198. Springer, 2007.
7. Björn Gottfried, Hans W. Guesgen, and Sebastian Hübner. Spatiotemporal Reasoning for Smart Homes. In Juan Carlos Augusto and Chris D. Nugent, (eds.), Designing Smart Homes, LNCS, vol. 4008, pp. 16–34. Springer, 2006.
8. Hua-Mao Gu, Xun Wang, Yun Ling, and Jin-Qin Shi. Building a Fuzzy Ontology of Edutainment Using OWL. In Yong Shi, G. Dick van Albada, Jack Dongarra, and Peter M. A. Sloot (eds.), ICCS 2007, LNCS, vol. 4489, pp. 591–594. Springer, 2007.

9. Hans Guesgen. Fuzzifying Spatial Relations. In Pascal Matsakis and Les M. Sztandera (eds). Applying Soft Computing in Defining Spatial Relations, pp. 1–16. Physica-Verlag GmbH, Heidelberg, Germany, 2002.
10. Hans Geusgen. Reasoning About Distance Based on Fuzzy Sets. Applied Intelligence, 17(3), pp. 265–270 (2002).
11. Jeffrey Hightower and Gaetano Borriello. Location Systems for Ubiquitous Computing. Computer, 34(8), pp.57–66 (2001)
12. Ken Hinckley, Jeff Pierce, Mike Sinclair, and Eric Horvitz. Sensing Techniques for Mobile Interaction. In Proc. of the 13th Annual ACM Symposium on User Interface Software and Technology (UIST 2000), pp. 91–100. ACM Press (2000)
13. Céline Hudelot, Jamal Atif, Isabelle Bloch. Fuzzy Spatial Relation Ontology for Image Interpretation. Fuzzy Sets and Systems, 159(15), pp.1929–1951, Elsevier (2008)
14. Gerd Kortuem, Christian Kray, and Hans-Werner Gellersen. Sensing and Visualizing Spatial Relations of Mobile Devices. In Proc. of the 18th Annual ACM Symposium on User Interface Software and Technology (UIST 2005), pp. 93–102, ACM Press (2005)
15. Krispin Leydon. Sensing the Position and Orientation of Hand-Held Objects: An Overview of Techniques. Technical Report, University of Limerick, December 2001.
16. Petteri Nurmi and Sourav Bhattacharya. Identifying Meaningful Places: The Non-Parametric Way. In Jadwiga Indulska, Donald J. Patterson, Tom Rodden, and Max Ott, (eds.), Pervasive 2008, LNCS, vol. 5013, pp. 111–127. Springer, 2008.
17. Ichiro Satoh. A Location Model for Pervasive Computing Environments. In Proc. of the 3rd Int. Conf. on Pervasive Computing and Communications (PerCom 2005), pp. 215-224. IEEE Computer Society, 2005.
18. Giorgos Stoilos, Giorgos Stamou. Extending Fuzzy Description Logics for the Semantic Web. In Christine Golbreich, Aditya Kalyanpur, and Bijan Parsia (eds). Proceedings of the Workshop on OWL: Experiences and Directions (OWLED 2007), Innsbruck, 2007.
19. Umberto Straccia. A Fuzzy Description Logic for the Semantic Web. In Elie Sanchez (ed.) Fuzzy Logic and the Semantic Web, Capturing Intelligence, chapter 4, pp. 73–90. Elsevier, 2006.
20. Geert Vanderhulst, Kris Luyten, and Karin Coninx. ReWiRe: Creating Interactive Pervasive Systems That Cope with Changing Environments by Rewiring. In Proc. of the 4th IET International Conference on Intelligent Environments (IE 2008), 2008.