To be or not to be Aware: Reducing Interruptions in Pervasive Awareness Systems

Sobah Abbas Petersen Jörg Cassens Anders Kofod-Petersen Monica Divitini
Department of Computer and Information Science
Norwegian University of Science and Technology
7094 Trondheim, Norway
sap,cassens,anderpe,divitini@idi.ntnu.no

Abstract

In this paper we look at awareness systems that use mobile and ambient devices for collecting or presenting awareness information and operate within an Ambient Intelligence Environment. Our focus is on Pervasive Awareness Systems (PAS) that mediate awareness with the aim to improve the feeling of being connected. In particular, we concentrate on the challenges that are connected to participation in multiple communities, requiring a continuous balancing between the need to "keep in touch" and to reduce interruptions. In the paper, we discuss software agents as a possible solution and identify the different roles that agents can play in reducing interruptions.

1. Introduction

As mobile and ubiquitous computing becomes a reality, awareness systems for keeping in touch are becoming more and more popular. Systems such as IM, Skype, Facebook are currently widely used to help the members of distributed social networks to feel connected. Lately, both industry and research have started exploring the possibility to fulfil the need to stay in touch using mobile and ambient devices [7, 13, 23]. In this paper, we focus on Pervasive Awareness Systems (PAS), i.e. awareness systems that use mobile and ambient devices for collecting or presenting awareness information and operate within an Ambient Intelligence (AmI) Environment.

The concept of AmI has been described as human beings surrounded by intelligent interfaces supported by computing and networking technology that is embedded in everyday objects such as furniture, clothes and the environment. The environment should be aware of the presence of a person (the user) and perceive the needs of the user and respond intelligently to these needs. We see AmI as an overlap of a number of paradigms; Ubiquitous Computing [29], Pervasive Computing [25] and Artificial Intelligence [24]. The Ubiquitous Computing aspect addresses the notion of accessibility of the technology. Pervasive Computing deals with architectural aspects and Artificial Intelligence techniques provide the context awareness and sensitivity to establish the user's needs and the appropriate response. Thus, PAS may collect awareness information from sensors in the environments and can present it to the user by using devices in the environment, e.g. ambient displays, lamps and mobile phones. These systems have the potential to improve people's connectedness by improving integration with everyday life activities.

Realising this potential raises some challenges related to the need to support participation in multiple communities. Most of the existing PAS focus on mediating awareness between two people. However, mediation of awareness always happens in a social context that gives meaning to it and relates to multiple social networks [16]. Empirical investigation has shown that we are continuously moving across different collaboration spheres that compete to get our attention [9]. While this study focuses mainly on the working environment, an initial study that we have conducted with a group of exchange students has pointed out similar issues also in connection to awareness mediation. Students struggle to keep in touch with their local as well as remote communities in a balance that is difficult to find. The information that can help us to keep connected to the members of one community can easily become an interruption to the participation to another, e.g. a call from the family during a work meeting. In this perspective, interruption is not only related to short term interruptions of ongoing activities, but rather the need to nurture social relationships that might last over a long period of time. Studies of interruptions and their computer support have mainly focused on the impact on individual tasks e.g. [4]. More recently, the importance to look at connected social issues has been recognised [10]. In

978-0-7695-3367-4/08 $25.00 © 2008 IEEE
DOI 10.1109/UBICOMM.2008.35

this paper we discuss the possibility to support users in the management of interruptions with software agents.

The work reported in this paper is a part of the ASTRA project[1]. The project goal is to create a service-oriented architecture and tools for enabling end-user communities to create their own PAS.

The paper is organised as following. Section 2 provides a brief state of the art. Section 3 sketches our perspective on PAS. Section 4 discusses interruptions and identifies two high level requirements. Sections 5 and 6 discuss how software agents could be used to meet these requirements.

2. Related Work

As ubiquitous computing becomes more a part of our everyday life, there is a challenge to coordinate the different devices in our environment and our interaction with them and the cognitive overload that this could lead to. An overview of research in this domain is available from Roda and Zajicek [21]. The authors have categorised the work in three broad areas: i) detection of the user's attention, ii) attention-aware information presentation, iii) support of shared and individual attention in collaborative situations. However, most of the work that was referenced in the article (as presented at the AMUCE 2007 workshop) was focused on attention management within a work context or collaborative work context, for example, by adapting the ways in which information is presented to the users.

A model for mediating interruptions was proposed by Dabbish and Baker where they consider the importance of the interruption and the interruption threshold [6]. They used the metaphor of an administrative assistant as interruption mediator. The interruption threshold takes into account the importance of the interrupter and the content of the interruption. In another paper, Dabbish and Kraut presented a study on interruption management within a team of collaborators, where awareness information is displayed and this information is used by the team members in deciding if they could interrupt their team members [5]. The awareness information displayed in this case reflected the team member's workload at any time. Tullio et al. [27] describe a display that could be used to support direct reports to a manager in deciding whether the manager can be interrupted or not. The display used the metaphor of ceramic tags on doors to indicate availability of people. The display showed a solid colour to indicate interruptibility and a gradient colour scale which ran from red (very uninterruptible) to green (very interruptible). The availability was based on a statistical model of interruptibility built by prompting the managers for a self-report of their interruptibility and sensors in the office detecting activity (e.g. typing on the keyboard). There has also been work on detecting an "appropriate" moment to interrupt a person by using physiological metrics such as the size of the pupil to indicate a person's workload [1]. Perhaps more relevant to the context of PAS is the use of interruption management in a social setting. Shell et al. [26] describe augmentation of computing devices with sensors, e.g. a mobile phone that can sense (via an augmented device) whether the user has eye contact with another person inferring that the user is having a conversation with someone. This information can be used by the mobile phone to determine if the user should be interrupted when there is an incoming call.

The work on interruption management focuses on situations within a work environment or collaborative work situation. The goal is to display information to a user so that she can use that information in deciding whether someone can be interrupted or not. This is different from a scenario where a user is able to decide how she wants to control whether she can be interrupted by specific persons or not. For example, there is a need to provide users tools to control the extent to which interruptions should reach them, depending on such things as what they are doing or who the interruption is from [2].

3. Pervasive Awareness Applications: Focus and Nimbus

In ASTRA, we look not only at individual users, but also at communities or social networks. Our starting point is that awareness is always mediated within communities. Even when the mediation process involves only two persons, awareness information gets meaning within the boundary of a larger community that defines e.g. expectations and conventions. Here we use the term community in a rather informal way to refer to "A number of individuals, defined by formal or informal criteria of membership, who share a feeling of unity or are bound together in relatively stable patterns of interaction" [28]. Awareness might be mediated within different types of communities, for example co-located and remote families, communities of friends, student communities, neighbourhoods. These communities vary in terms of dimensions, accepted forms of participation, needs and bonds. What is acceptable in one community may be unacceptable in another. In general, each individual can be a member of multiple communities. Within each community, diverse rules and conventions might influence how issues such as awareness service provisioning, sharing and tailoring should be facilitated by the system.

It is also important to note that in PAS, collection and fruition of awareness information is situated in specific physical spaces which have different characteristics in terms of the level of support that can be provided as well as of rules and conventions. More important than the actual capa-

[1] http://www.astra-project.net/

bilities of the specific space is how a community inhabits it, possibly sharing it with other communities. This takes us to the critical distinction between space and place [11]. While space indicates the physical surroundings, the place indicates social constructs to which people associate a meaning, norms of behaviour, and the like. A typical example to illustrate this distinction is the notion of a house, which is a place, and a home, i.e. a space that is turned into a place by the social and emotional meaning that we build around it.

Following is a simple scenario of PAS in ASTRA (see also the video demonstration available at the project web site.)

Alice is in her office and wants to go for a walk to take a break. She can use an ASTRA enabled device, for example a cube with motion and orientations sensors, to express her wish for a walk by giving this cube a specific orientation. She has previously decided to make this information available to her friends, who can individually define how this information is presented. When Bob is at home, he has chosen to see updates of Alice's state on his digital picture frame, showing a picture of Alice. Bob can then decide whether he wants to contact Alice for a walk. When he is on the move, he has set the system so that any notification from Alice is sent as a SMS on his mobile phone.

In ASTRA, for the time being, both publishing of awareness states and subscription of other person's states are defined via user defined rule sets. In our example, Alice has coupled the orientation of her cube with the wish for walk, and Bob has defined in his rule set that his subscription of Alice's wish for a walk should be displayed on his picture frame or, depending on his location, as an SMS.

The theoretical background for ASTRA awareness applications can be found is the focus-nimbus model, originally described by Benford et al. [3]. The authors use room metaphors as the basis for a spatial model to support communication between participants in virtual rooms. The basic idea is that people can not only visit different virtual rooms, but they can move around in these different rooms, and the (modelled) spatial characteristics of the rooms mediate the communication between different persons in the room. Two concepts are introduced; the focus represents a space in the room where a person targets his attention. People are more aware of objects in the focus than those outside. The nimbus is the counterpart, representing where the person locates himself in the room. Objects are more aware of a person if the object is located in the person's nimbus than when it is located outside [11, p. 220]. Awareness is defined through the interaction of focus and nimbus, and can be mathematically expressed through the spatial relation of a focus and a nimbus. This model has been generalised by Rodden [22]. He extends the notions of focus and nimbus towards application areas without an explicit notion of spatial relations. Basically, he introduces a graph model for a domain, and awareness becomes a property of this graph. In its simplest form, the awareness measure is the length of the path between two users. Metaxas and Markopoulos have later presented a formal model which concentrates on the communication aspects of the focus-nimbus model [18]. Their model addresses issues of privacy by allowing for plausible deniability and deception. In the terms of our example, Alice would make her wish for walk available by placing it on her nimbus. She can control to which community of users she publishes this aspect. Bob, on the other hand, would have a focus on this particular aspect of Alice's information.

4. Interruptions in Pervasive Awareness

When we discuss interruptions in the context of PAS, we have to take into account the fact that the need to prevent interruptions might be short term, e.g. when someone is talking on the phone or in a meeting; or relate to a longer period, e.g. a student reducing social activities during the exam period and deciding to be in contact only with the closest friends.

Interruptions in PAS might be managed in two ways: by acting of one's nimbus and by acting on one's focus. Under a condition X (in which interruptions are not welcome) the user might decide to publish some information that will make other people aware of her current status, e.g. "I am at work". In this case the responsibility is left to the person receiving the notification and relies on the existence of a set of shared conventions, e.g. "when someone is at work, they should not be disturbed unless. . . "

Traditional applications often use only a limited set of states that do not convey context information and therefore do not support the emergence of these conventions. In general, we could assume that the more awareness is provided to others about one's own status, the more others will be able to perform accordingly, and, in this case, reduce the number of unwanted interruptions. The advantage here is that the same information can be published once to multiple communities, leading to completely different behaviours based on the conventions of the community. For example, "I am at work" might trigger completely different behaviours by one's family and by one's colleagues.

Acting on his focus, under a condition X (in which interruptions are not welcome) the user might decide to block the notification of some information about other people status, for example "Alice wants to go for a walk." In this case interruptions are blocked before they might happen.

Acting on focus and nimbus can be done by defining rules specifying when certain information should be conveyed to the members of a community and when certain information should not be conveyed or visualised. Defining these rules is time demanding and might go far beyond the capabilities of normal users. Therefore we require:

Requirement 1: *PAS should support users in the definition of their focus and nimbus taking into account the need to reduce interruptions.*

For example, rather than defining rules for each specific application, the PAS could help the user to define a filter based on his location (when I am at home, then visualise application X [Alice wants a walk] by displaying Alice's picture or do not visualise Y [Alice is at home].) The system could also suggest rules based on the behaviour of other users in the same location.

We have however to accept that it is impossible to foresee all the possible situations that might emerge and the ways that mediation of awareness might be perceived as disruptive. Therefore we add:

Requirement 2: *PAS should provide support for management of events that might be perceived by users as interruptions.*

For example, depending on the situation, the system could decide to visualise a message in a modality that is less disruptive (e.g. a SMS rather than a phone call).

Agents and multiagent systems have been suggested as an important paradigm in the design of intelligent environments [19]. Agents are specially good at modelling real world and social systems due to their social abilities (agents, by definition, are social entities [30]). We propose agents as a suitable solution technology to meet the two requirements stated above in achieving interruption management in PAS. In the next sections, we discuss how software agents may be used to address the requirements.

5. User Definition of Focus and Nimbus

The rules that are used can be simple ones that take into account one parameter such as the time of the day, e.g. I do not want any interruptions from anyone, from 01:00-06:00hrs, everyday. Here, the user sets a time as the context and it applies to everyone that she interacts with. Such a rule is simple from a programming perspective. However, it does not reflect the reality of a person that has many social relationships. For example, if Bob is working at home, he may not want to be disturbed by anyone except his closest colleagues who are cooperating with him on a task. In the first example of the rule, by stating a rule explicitly, the user is also implicitly telling the system that he could be disturbed by everyone he has contact with any other time.

To find the appropriate balance of interruptions desired by a user at any time, the types of rules and the amount of rules that are required may be quite complex and take up a lot of the user's time. It may also require duplication of rules and ensuring that one rule does not conflict another. A review of the literature in the use of agents in such situations suggests that agents can play three different types of roles to assist the user in defining his interruption management rules: (1) personal assistants, (2) advisors and (3) recommenders.

5.1. Agents as Personal Assistants

Agents could support the user in defining rules by playing the role of a personal assistant, where the agent performs most of the tasks and relieves the user of the work. Maes proposed a means of complementing direct manipulation with indirect manipulation, where the user collaborated with an agent to achieve a task. The metaphor of a personal assistant was suggested to illustrate a series of agents that supported a user in managing the work and information overload [15]. For example, a user could delegate the effective management of email or the responsibility of scheduling a meeting to an agent. In these examples, the user specified some information to the agent that the agent could use to manage the specified task.

To avoid the user having to provide all the rules or the information to create the rules explicitly, i.e. the competence of the agent, algorithms from machine learning have been applied. The mental models formed by the agents about their users must be trustable by their users [27]. So, it is important that the agents acquire knowledge about the user that represents the user's behaviour and desires. In addition to traditional machine learning techniques such as programming by example, a user can teach or increase an agent's competence in several ways: by providing direct feedback to the agent based on its actions and by providing explicit examples to illustrate how to behave in certain situations. In addition to this, agents should also be designed such that they can ask for advice from the user or other agents when they encounter new situations. This capability of agents is particularly useful in PAS. For example, members of a community may have similar desires in how they would like to manage their interruptions. Thus, providing information about user's rules via their agents could be useful for newcomers to a community or agents that are looking for advice. Agents can build their competence based on the collective competence of the agents that belong to the community of users.

5.2. Agents as Advisors

Another role that an agent could play in assisting a user in defining his focus and nimbus is that of an advisor [14]. Here, the user is the one that defines the rules, not the agent. The agent teaches or advises the user about how to define the rules. According to the authors, advisory agents have shown experimental evidence that they improved user performance.

Advisory agents are definitely helpful and necessary in ensuring that users can define their focus and nimbus effectively. Depending on the user interface that is provided by the system, certain groups of users, e.g. non-technical users or the elderly, may require advice in defining the rules. Also, if a diverse range of user interfaces is provided by the system, which may be desirable from the perspective of catering for different types of users, it may also require more from the users in learning how to define rules using the different interfaces.

5.3. Agents as Recommenders

We also see a need for agents that can recommend rules to users. An agent may know about the rules that other users have used for a specific application, within a specific community and this may be helpful for the user to know.

Recommender systems have been popular in the field of electronic commerce and companies such as Amazon use ideas from recommender system to suggest products that they believe may be of interest to their users. In particular, due to the demand of personalisation technologies, there has been an increase in the popularity of recommender systems [8]. Recommender systems fill the role of the social process of recommendations by people either by word of mouth, book or movie reviews [20]. There are two main techniques that are used by recommender systems: collaborative filtering and content-based filtering. Collaborative filtering is based on the behaviour of similar people, e.g. a rule used by Bob in defining his interruption management when he's working from home may be recommended to his colleagues as they may have similar profiles. Content-based filtering is based on features of items that a user liked to recommend items in the future. For example, an awareness application that is used by Alice, such as the wish for a walk, could be recommended to other users who have expressed interest in outdoor activities.

The roles played by all three types of agents seem similar. However, these are distinct, yet complementary roles that we believe are necessary to support users in defining their rules.

6. Interruption Management

There have been attempts to develop models to support reasoning about information awareness vs. interruption of the user [12]. In addition to defining the rules for managing interruptions, the user also needs to be able to manage her interruptions in a context-sensitive and effective manner. One of the interaction modes of a user could be via his personal agent that decides whether the user could be interrupted [6].

McFarlane proposed a taxonomy of human interruptions where he identified four ways in which people manage the interruptions they receive: (1) immediate interruption; (2) negotiated interruption; (3) mediated interruption; and (4) scheduled interruption (or coordination by prearranged convention or explicit agreement) [17]. For example, some situations may require that the agent conducts an explicit or implicit negotiation with the user before a message is passed onto the user. Similarly, the agent may be a mediator for the user. For example, when Alice is in a meeting at work, she may assign an agent to receive messages from her family which are conveyed to her when she is taking a break from her revision. This ensures that she does not miss any messages from her family, while she is preserving her interruption management model. She could also assign her agent to convey some sort of message to her family to let them know that she will respond in due course, so that the sender of the message is not offended or worries unnecessarily about Alice's lack of response. One advantage in using agents for managing interruptions is that agents are proactive by nature and they could find an appropriate moment to convey the message to Alice rather than waiting for Alice to always initiate it. Such agents could also be used for preserving the social intelligence of the system.

7. Conclusion

As ubiquitous computing becomes a reality and a part of our everyday life, there is a challenge to manage our awareness of social connections and interruptions. In this paper, we have discussed the problem of interruption in PAS and we have identified different roles that software agents can play in supporting it. The paper does not aim to provide a simple answer to the problem, but it addresses its complexity and identifies challenges to PAS. At present, we are continuing to use the focus and nimbus model in the context of multiple communities. In the future, we plan to design the architectures for the different types of the agents proposed and work towards integrating them into the service-oriented architecture of the ASTRA system.

Independent of the role played by agents, they should be able to support users (1) in acting according to the rules of conducts of the community within which the interaction is happening; (2) in acting according to the rules of conduct of the place where the interaction is happening. In this perspective, social intelligence emerges as a central design guideline for PAS. In doing so, we should look at two complementary aspects. The systems should help users to behave in a socially intelligent way, or at least it should not disrupt their capability to do it. The systems should also be perceived as socially intelligent.

8. Acknowledgement

The research reported in this paper is financed by the EU project ASTRA-EU029266. Parts of this work has been

supported by Accenture Innovation Lab Norway.

References

[1] P. D. Adamczyk, S. T. Igbal, and B. P. Bailey. A method, system, and tools for intelligent interruption management. In M. Sikorski, editor, *Proceedings of the 4th International Workshop on Task Models and Diagrams*, pages 123–126. ACM Press, 2005.

[2] P. Andrews. Vying for your attention: Interruption management. *Executive Technology Report*, (7):1–8, 2007.

[3] S. Benford, A. Bullock, N. Cook, P. Harvey, R. Ingram, and O.-K. Lee. From rooms to cyberspace: models of interaction in large virtual computer spaces. *Interacting with Computers*, 5(2):217–237, 1993.

[4] M. Czerwinski, E. Horvitz, and S. Wilhite. A diary study of task switching and interruptions. In *Proceedings of ACM CHI 2004 Conference on Human Factors in Computing Systems*, pages 175–182, 2004.

[5] L. Dabbish and R. Kraut. Controlling interruptions: Awareness displays and social motivation for coordination. In J. Herbsleb and G. Olson, editors, *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 182–191. ACM Press, 2004.

[6] L. A. Dabbish and R. S. Baker. Administrative assistants as interruption mediators. In G. Conckton and P. Korhonen, editors, *CHI '03 extended abstracts on Human factors in computing systems*, pages 1020–1021. ACM Press, 2003.

[7] A. K. Dey and E. de Guzman. From awareness to connectedness: the design and deployment of presence displays. In R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, editors, *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 899–908. ACM Press, 2006.

[8] A. Felfernig, G. Friedrich, and L. Schmidt-Thieme. Recommender systems. *IEEE Intelligent Systems*, 22(3):18–21, 2007.

[9] V. M. González and G. Mark. Managing currents of work: multi-tasking among multiple collaborations. In H. Gellersen, K. Schmidt, M. Beaudouin-Lafon, and W. Mackay, editors, *Proceedings of the Ninth European Conference on Computer Supported Cooperative Work*, pages 143–162. Springer, 2005.

[10] R. Harr and V. Kaptelinin. Unpacking the social dimension of external interruptions. In T. Gross and K. Inkpen, editors, *Proceedings of the 2007 International ACM Conference on Supporting Group Work*, pages 399–480. ACM Press, 2007.

[11] S. Harrison and P. Dourish. Re-place-ing space: The roles of place and space in collaborative system. In G. Olson, J. Olson, and M. S. Ackerman, editors, *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative work*, pages 67–76. ACM Press, 1996.

[12] E. Horvitz, C. Kadie, T. Paek, and D. Hovel. Models of attention in computing and communications: From principles to applications. *Communications of the ACM*, 46(3):52–59, March 2003.

[13] I. Keller, W. van der Hoog, and P. J. Stappers. Gust of me: Reconnecting mother and son. *IEEE Pervasive Computing*, 3(1):22–27, 2004.

[14] H. Lieberman and T. Selker. Agents for the user interface. In J. Bradshaw, editor, *Handbook og Agent Technology*. AAAI/MIT Press, 2000.

[15] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, 1994.

[16] G. McEwan and S. Greenberg. Supporting social worlds with the community bar. In K. Schmidt, M. Pendergast, M. Ackerman, and G. Mark, editors, *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, pages 21–30. ACM Press, 2005.

[17] D. C. McFarlane. *Interruptions of People in Human-Computer Interaction: A General Unifying Definition of Human Interruption and Taxonomy*. PhD thesis, Human-Computer Interaction Research Group Department of Computer Science George Washington University, 1997.

[18] G. Metaxas and P. Markopoulos. 'aware of what?' a formal model of awareness systems that extends the focus-nimbus model. In *Proceedings of the IFIP conference EHCI 2007*. Springer, 2007.

[19] C. Ramos, J. C. Augusto, and D. Shapiro. Ambient intelligence – the next step for artificial intelligence. *IEEE Intelligent Systems*, 23(2):15–18, 2008.

[20] P. Resnick and H. R. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.

[21] C. Roda and M. Zajicek. Attention management in ubiquitous computing environments (AMUCE 2007). Workshop at UBICOMP 2007, 2007.

[22] T. Rodden. Populating the application: A model of awareness for cooperative applications. In *CSCW '96: Proceedings of the 1996 ACM conference on Computer Supported Cooperative Work*, pages 87–96, New York, NY, USA, 1996. ACM Press.

[23] N. Romero, P. Markopoulos, J. van Baren, B. de Ruyter, W. Ijsselsteijn, and B. Farshchian. Connecting the family with awareness systems. *Personal and Ubiquitous Computing*, 11(4):299–312, 2006.

[24] S. J. Russel and P. Norvig. *Artificial Intellignece: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition, 2002.

[25] M. Satyanarayanan. Pervasive computing: Vision and challenges. *IEEE Personal Communications*, 8(4):10–17, August 2001.

[26] J. S. Shell, T. Selker, and R. Vertegaal. Interacting with groups of computers. *Communications of the ACM*, 46(3):40–46, 2003.

[27] J. Tullio, A. K. Dey, J. Chalecki, and J. Fogarty. How it works: A field study of non-technical users interacting with an intelligent system. In M. B. Rosson and D. J. Gilmore, editors, *CHI*, pages 31–40. ACM, 2007.

[28] B. Wang, J. Bodily, and S. K. S. Gupta. Supporting persistent social groups in ubiquitous computing environments using context-aware ephemeral group service. In *Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications (PerCom'04)*, pages 287–296. IEEE, 2004.

[29] M. Weiser. The computer for the 21st century. *Scientific American*, pages 94–104, September 1991.

[30] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.