Where the Footprints Lead: Tracking Down Other Roles for Social Navigation

Paul Dourish
Xerox Palo Alto Research Center, Palo Alto, California
dourish@parc.xerox.com

Abstract

Collaborative Filtering was proposed in the early 1990's as a means of managing access to large information spaces by capturing and exploiting aspects of the experiences of previous users of the same information. Social navigation is a more general form of this style of interaction, and with the widening scope of the Internet as an information provider, systems of this sort have rapidly moved from early research prototypes to deployed services in everyday use.

On the other hand, to most of the HCI community, the term "social navigation" is largely synonymous with "recommendation systems": systems that match your interests to those of others and, on that basis, provide recommendations about such things as music, books, articles and films that you might enjoy. The challenge for social navigation, as an area of research and development endeavour, is to move beyond this rather limited view of the role of social navigation; and to do this, we must try to take a broader view of both our remit and our opportunities.

This chapter will revisit the original motivations, and chart something of the path that recent developments have taken. Based on reflections on the original concerns that motivated research into social navigation, it will explore some new avenues of research. In particular, it will focus on two. The first is social navigation within the framework of "awareness" provisions in collaborative systems generally; and the second is the relationship of social navigation systems to spatial models and the ideas of "space" and "place" in collaborative settings.

By exploring these two ideas, two related goals can be achieved. The first is to draw attention to ways in which current research into social navigation can be made relevant to other areas of research endeavour; and the second is to re-motivate the idea of "social navigation" as a fundamental model for collaboration in information-seeking.

1 Introduction

In the early 1980s, researchers working in the area of interactive systems became increasingly interested in the topic of cooperative work. Human beings are, after all, social animals, and most activities in which we engage are conditioned by and conducted in coordination with other individuals. We work collectively. However, until this point, the focus of interactive systems research had largely been on a single user sitting at a desk in front of a computer screen. The field of Computer-Supported Cooperative Work (CSCW) emerged in response to this focus. It drew attention to the range of concerns that lay outside this computer/human dyad but which fundamentally affected the nature of work at the computer, such as the social setting in which the activity took place, and the role that that the activity played in an individual's collaborative actions. In the years that have followed, CSCW has legitimised this concern and reoriented HCI to take into account the social context in which work is conducted. At the same time, it has also had an influence on the development of everyday technologies for computer-based work; as the Internet has become a more commonplace computational phenomenon, so technologies such as

workflow and groupware have moved out of the research laboratory and into everyday computational practice.

The topic of social navigation offers an opportunity to take this even further. Social navigation is one of the most direct expressions of the fundamental principle that the action of a user at a computer is driven not simply by a set of internal goals and cognitive processes, but by the social setting in which they find themselves, by the action of others, individually and collectively, and by the social nature of the work being conducted and the goals sought. In other words, social navigation provides an opportunity to take those aspects of computer-based interaction and information seeking which might still be regarded as primarily single-user activities, and to invest them with a sense of the social. In systems supporting social navigation, the social and collaborative aspects of work are not limited to the overtly cooperative activities such as meeting support, communication and shared workspace manipulation, but to activities that would seem, still, to be grounded in the model of a single user sitting at a computer. Reading on-line documents, navigating file systems, searching a database, finding a book to read and sorting electronic mail might all seem like fundamentally individual activities; private and encapsulated in one person's head. However, the use of social navigation techniques allows us to be able to capitalise upon the social nature of the everyday world, and so to enrich the interface with collaborative support for individual tasks.

This volume presents a wide variety of current research into the role and function of social navigation. This chapter is concerned with potential new opportunities for the application of the social navigation concept in interactive systems in general, and will argue that social navigation is a more general phenomenon than current practice would suggest. As a result, the focus here will be less on current research on social navigation and more on those other areas to which it might be applied. The body of the chapter will discuss how we might apply understandings of social navigation to address wider problems in the design, development and analysis of interactive and collaborative systems. A consideration of new opportunities in research on social navigation, however, should begin by looking at the origins and scope of the term, and so that will be the first topic.

1.1 Perspectives on Social Navigation

How should we consider social navigation? What range of phenomena and systems does the term name, and how do we come to consider them as aspects of a single concept?

Although there had been investigations of collaborative systems as far back as the NLS/Augment work of Engelbart and his colleagues in the 1960's [17] and which persisted throughout subsequent developments in data communication networks, it was only in the early 1980's that the term "Computer-Supported Cooperative Work" was coined and the field began to develop a more uniform identity and a shared base of understandings from which to build. At this point, a variety of systematic investigations began into the role of mutual cooperative behaviour through computer systems, the role of computational technologies to support coordination and collaboration, and the use of data-sharing to support concerted action.

The Tapestry system [22, 43], developed at Xerox PARC, introduced the idea of "collaborative filtering", which became one of the key ideas that has driven subsequent work on social navigation. Collaborative filtering systems provide a user with recommendations of their likely interest in data items on the basis of "interest matches" derived from previously ratings from a set of users. Tapestry operated over a database of electronic mail messages and allowed users to exploit "votes" that previous users had entered concerning the value or interest of the message. By selecting an appropriate group of users based on overlaps in interest, Tapestry could allow an individual user to exploit collective experience.

Recommender systems, as these sort of applications have become known, have become one of the most visible forms of social navigation in everyday systems. The work of Pattie Maes' group at the MIT Media Lab [39], Will Hill and his colleagues at Bellcore [29] and others has explored the use of collaborative recommendation of books, films and music, and these technologies have been incorporated into deployed technologies such as electronic commerce systems. Recommendation systems vary along a number of different dimensions, but the basic insight is similar to that of the collaborative filtering system in that they provide users with a means to use information that others have left behind as a cue to exploring an information space. In particular, recommendation systems typically use information about individual interests (e.g. which books they like and which they dislike) to build a profile which can be matched with the profiles of other individuals. If the system can determine that two people have similar tastes in books (because there is enough overlap in their stated likes and dislikes to imply a correspondence), then it can suggest that books that one likes but the other hasn't seen are probably good ones to recommend. In electronic commerce settings (such as an on-line store), profiles can be developed on the basis of shopping habits, reducing the overhead in generating the profiles in the first place.

1.1.1 Spatial, Semantic and Social Navigation

In a short paper presented at HCI'94, Dourish and Chalmers [16] introduced the term "social navigation", which included systems of this sort. However, the focus of that brief exploration was broader, and the term "social navigation" was coined to describe a particular phenomenon, in which a user's navigation through an information space was primarily guided and structured by the activities of others within that space.

"Social" nagivation was contrasted with "spatial" and "semantic" navigation. Spatial navigation relies on the structure of the space itself, often a two- or three-dimensional metaphor of some spatially-organised real-world phenomenon (such as an office, street or landscape). Virtual reality systems, for example, place considerable reliance on spatial navigation, offering users a spatial organisation by which to explore an environment. "Semantic" navigation, in contrast, relies on the semantic structure of the space. A hypertext system, for example, provides "links" between semantically-related items and offers a means to move from one item to another according to these semantic relationships.

Dourish and Chalmers were attempting to clarify the distinctions between the three forms to support better evaluation of the important features of navigational systems. This concern arose from the emergence of collaborative virtual reality systems or "Collaborative Virtual Environments" (CVEs), with their strong appeal to spatial models. CVEs use a 3-D visualisation metaphor to organise data and interaction, and so are firmly grounded on a spatial model; by exploiting our familiarity with the spatial structure of the everyday world, they can provide smooth and intuitive interaction with large information spaces. In CVEs such as DIVE [9] and MASSIVE [23], a user may have visual access not only to data items or objects in the space but also to other individuals and to their interactions. This gives them the opportunity not only to interact with data items, but also to be able to see (and hence be influenced by) other people's interactions in the space; and researchers studying CVEs have developed a range of mechanisms for providing participants with a control over views of each other's actions [5][1]. However, the arguments that apply to the familiarity of interaction with objects in a three-dimensional space to do not apply unproblematically to interactions between individuals. Is inter-personal interaction in the three-dimensional space driven primarily by the structure of that space? Is spatial structure enabling social action, or is social action happening within a spatial structure?

---

[1] However, the restricted field of view and cumbersome mechanisms of navigation and exploration afforded by current CVEs can introduce problems for mutual orientation and reference [30].

This confusion arises because semantic and social navigation may well take place in spatially-organised settings. Semantic and social navigation do not name types of systems; rather, they name *phenomena of interaction*. They may all occur in the same sort of "space". For instance, imagine browsing in a bookstore. If I pick up a new book because it is sitting on the shelf next to one I've just been examining, then I'm navigating spatially. If I pick up another book because it was referred to in a citation in the first book, then I'm navigating semantically; and if I pick up yet another because it was recommended to me by someone whose opinion I trust, then I am navigating socially. The conceptual separation between "spatial", "semantic" and "social" styles of information navigation was intended to provide terms in which these different forms of behaviour could be discussed.

This is the basic principle from which the investigations to be presented here proceed: that social navigation is an *interactive phenomenon* rather than a class of technology. In this chapter, I will take the term "navigation" as an adequate term for all information-seeking activities. The space to be navigated is the notional space of potential information items, be those book recommendations, web pages, interesting data items, other people, or whatever. Navigation does not, in itself, imply "movement" in any typical, spatial sense. What will be of concern is how the information-seeking activities can be framed socially and how this framing can assist an individual in the course of their information-seeking activity.

1.1.2 New Opportunities for Social Navigation

In light of these considerations, we can take a number of different perspectives on social navigation. One would be to think of social navigation as an explicit activity, in which a user calls upon others to request advice or pointers, either directly or via an intermediate (perhaps an artifact like a FAQ list). Another, more commonly accepted amongst researchers in the field, is to think of social navigation as essentially an intrinsic element of individual interaction, one which respects that individual action is carried out within a complex of social relationships and provides a means to exploit this connectivity to other individuals to help people organise information.

In this chapter, some alternative conceptions of social navigation will be presented and explored. In the first, social navigation will be considered as an aspect of collaborative work, in which information can be shared within a group to help each group member work effectively, exploiting overlap in concerns and activities for mutual coordination. In the second, it will be presented as a way of moving through an information space and exploiting the activities and orientations of others in that space as a way of managing one's own spatial activity.

These last two perspectives will be the main topics for the rest of this chapter. The following section consider the notion of social navigation in the context of collaborative systems, and particularly with relation to the idea of "awareness" in collaborative workspaces. The subsequent section will consider the relationship between social navigation and the conceptions of "space" prevalent in collaborative and interactive systems. By taking these two particular perspectives, I hope to be able, first, to cast some light on the fundamental nature of social navigation broadly construed and, second, to open up new avenues for the use and design of social navigation facilities in interactive tools.

2 Social Navigation and Collaborative Awareness

The idea of social navigation is firmly based on the fact that information about others and about others' activities can be beneficial to an individual in the conduct of their activity. In collaborative systems, this idea is clearly strongly related to the topic of the mutual "awareness" between collaborators of each other's activity. In this section, I will pursue this relationship in more detail and use the perspective of social navigation to highlight opportunities for a broader conception of collaborative awareness.

2.1 Awareness in Collaboration

The topic of "awareness" has become a significant focus of research activity in the CSCW domain. The principal idea behind this notion is that the activity of others within a group setting can be a critical resource for the individual in managing their own work [14]. The term "awareness" has come to refer to the ways in which systems can present this information and so provide to participants in a collaborative activity an "awareness" of the action and activity of others.

This feature of collaborative activity is clearly not restricted to electronically mediated settings. Excellent examples can be found, for instance, in the work of Heath and his colleagues studying a variety of collaborative settings such as transit system control rooms [26] and dealing rooms [27]. In these settings, a common observation is that the detail of coordination between individuals is tied to a "peripheral monitoring" by which they maintain an understanding of the state of each other's activities, so as to be able to organise their own behaviour in correspondence. At the same time, individuals also choose to organise their work so that it can be successfully monitored; that is, they display the state of their own activity.

Consider an example. In their study of a London Underground control room, Heath and Luff [26] present a detailed analysis of an interaction between two individuals in the control room. One, the Line Controller (Controller), has responsibility for the movement of trains on the line, while the other, the Divisional Information Assistant (DIA), has responsibility for information provided to passengers about the service. Health and Luff observe that the DIA's actions can be seen to be tied to the precise detail of what the Controller is doing, to the extent that, on overhearing the Controller discussing a problem with a driver over the in-cab intercom, the DIA can tell that this problem is going to result in a service disruption and is *already beginning* to announce to passengers the consequences of this disruption *before* the Controller has informed him of the nature of the problem. This sort of very tightly coupled interaction depends upon a detailed understanding, on the part of the DIA, of the Controller's activity and the consequences it might hold. Conversely, when "reworking" the schedule (to accommodate changes in the service resulting from local disruptions), the Controller will talk through the changes he is introducing in such a way that he can be overheard by the DIA precisely to enable this sort of interaction.

In other words, what we find in studies of naturalistic collaborative work is that the twin mechanisms of peripheral monitoring and explicit "display" provide collaborators with the means by which to couple their work. These observations have motivated the development of technological approaches that provide for the framing of individual activity within the context of the group.

2.2 Designing for Awareness

The developers of CSCW systems have drawn from these observations the importance of awareness information for coordinating collaborative activity and so they have developed a range of mechanisms for providing awareness information in collaborative technologies, with considerable success in a variety of domains [42, 12, 21, 38]. However, the emphasis in these technical developments has largely been on the first of the two mechanism found in real-world settings, that is, on the peripheral monitoring. Systems have provided a variety of means for what Dourish and Bellotti [14] referred to as "shared feedback", by which collaborators can see the results of each others' actions. For instance, the SASSE system [2] provided a variety of ways by which users could see each other's work, including "radar views" and distortion effects on the display of the textual workspace. Gutwin et al. [24] provide an overview of a variety of techniques for visualising the activity of others.

These systems have provided for ways to see the work of others, and so have supported the "peripheral monitoring" aspect of the awareness-based coordination of activities that Heath and Luff have so vividly demonstrated in their studies. Relatively little work, however, has gone into the area of the explicit production of activity information for coordination purposes. In fact, in general, this question of how information is to be produced for awareness purposes – corresponding, for example, with Heath and Luff's observations of the mumbling Line Controller – has remained relatively unexplored so far.

Certainly, it raises some hard problems. It immediately leads to a set of concerns about what kinds of uses people expect the information to be put to, not to mention the sorts of privacy and protection issues raised, in the context of ubiquitous computing environments, by Bellotti and Sellen [1993]. However, some potential solutions to these problems lie in the ways in which similar issues have been tackled in research into social navigation.

2.3 Taking the Social Navigation Perspective

The fundamental relationship between social navigation and collaborative awareness technologies is simple, and lies in their mutual concern with the means by which an individual can exploit the information of others. What is interesting is that so many social navigation systems, and particularly the recommendation systems that have become so popular, are based not simply on the information of others, but information about the activity of others, which is fundamental to collaborative awareness. So I can capitalise upon information about what books other people have bought, what web pages they have visited, what newsgroup articles they have read, and so forth. The reason this tends to be the case is that activity information is easy to capture at low overhead to the user; it is much more convenient for my system to record which books you have bought from an online store than it is to present you with a list of three hundred books and ask you to rate them each on a scale from one to ten. However, despite this pragmatic explanation, the relationship between the two concerns is still suggestive.

It is particularly suggestive because it opens up opportunities for further cross-fertilisation. In particular, there are a number of features from social navigation systems that could fruitfully applied in systems providing collaborative awareness.

2.3.1 Aggregation

One interesting feature of social navigation systems in comparison to traditional collaborative awareness technologies is that of aggregation. By this, I mean that social navigation systems might bring together information from a number of individuals at the same moment, presenting either information about a plurality of users or about a fictional composite user.

Consider the case of the book recommendation system again. The sorts of observations we might expect from such a system are ones like, "People who liked 'The Thought Gang' also liked 'Mr. Vertigo', 'My Idea of Fun' and 'Towing Jehovah'". What is significant about this, from the perspective of collaborative awareness systems, it is that the information is not about a single individual, but rather, reports a trend which is based on observation from a variety of individuals. In all likelihood, there are some number of individuals who purchased all four books, but there may not be; what is presented is an aggregate observation based on data gathered from a number of individuals.

In collaborative awareness systems, however, the focus has normally been on presenting information from different individuals separately. Most displays tends to present me with the information about the current activity of Joe, Katie and Brian independently, presenting a fairly direct and faithful rendition of their current activity. Awareness systems generally do not provide any sort of interpolation or generalisation of information. And yet, at the same time, this

introduces serious problems about scalability. Since there is no generalisation, and only individual reports, problems arise over the number of individual reports that can be accommodated, and how they can be presented together. The social navigation approach offers an alternative that could be of considerable benefit in trying to address this problem.

2.3.2 Presentation

A related issue concerns the terms in which information is delivered and presented. Partially, perhaps, as a result of the potential aggregation of information from different sources, social navigation systems have had to consider the possibility that the terms in which information about the activity of others is presented to a user might not be the same terms in which it was gathered.

In the case of inherently aggregated information, such as votes, this is inevitable; the gathering of information is in terms of positive or negative votes about some item, whereas the presentation of the information is in terms of the mass of votes for or against; a quantitative measure, and so of quite a different sort. However, there are other ways in which the form of the information is altered for presentation – web pages rather than links, books rather than purchases, or whatever. In other words, the designers of social navigation systems, since they are focussed primarily on the notion of helping the end-user rather than simply presenting the information, are open to the fact that the form in which the information is most usefully presented might not be the same as that in which it was gathered; while the designers of collaborative awareness systems have often been concerned with verisimilitude as a primary design criterion, showing the movement of mouse cursors and text edit points as accurately as possible, or giving direct views of the activity of others (as in video-based systems such as Portholes [15]). Exploiting, again, the clear foundational relationship between social navigation and collaborative awareness opens this up to some question, and points the way towards new approaches to awareness that consider more how the information is to be exploited, and hence how it might best be transformed and presented.

2.3.3 Decoupling

One final feature of social navigation systems relevant here is that they are both synchronous and asynchronous. Most awareness systems, on the other hand, tend to be synchronous only, displaying the user information about other collaborating individuals in real-time, but giving them less information about what the others may have done in the past when the user was disconnected or engaged in another task. Although there has been some work on asynchronous means for conveying awareness information, such as Hill et al.'s "edit wear" and "read wear" [28], collaborative awareness systems have largely focussed on synchronous solutions. The value of social navigation systems as an asynchronous mechanism for sharing activity information and for capitalising upon such shared information is, surely, ample indication of the value to be derived from asynchronous collaborative awareness.

The problem here is one of coupling. In the typical, synchronous, collaborative case, the presentation of information about a user's activity is directly coupled to the activity itself. As surely as one user's keystrokes result in characters being entered into a text buffer, they also result in indications of typing activity on another user's screen, and so on. There are not, typically, mechanisms for decoupling the awareness information (although some awareness "widgets" are a start in that direction [1]). In the same way that technologies for social navigation aggregate and re-represent information about individual actions for presentation as indicates of group trends, they also introduce a decoupling that extends the usefulness of this information beyond the synchronous case.

2.4 Summary

What this exploration suggests, then, is that the concerns of collaboration awareness and the phenomenon of social navigation are closely related. They both rely on being able to present information about the activity of individuals in a way that allows other individuals to capitalise on it for the management of their own activities. What is more, by looking at collaborative awareness through the lens of social navigation, we have been able to identify a number of aspects of system design – aggregation, presentation and decoupling – that could help address current problems in the development of awareness technologies when we think of awareness as being a form of social navigation and information management.

The phenomenon of social navigation can also be applied to other areas of concern in collaborative systems. In particular, we will now go on to explore it in another context. Taking the term "navigation" more seriously, we can consider social navigation in relation to recent trends investigating the notions of "space" and "place" in CSCW systems.

3 Social Navigation and Models of Space

The very term "social *navigation*" invokes a spatially-organised world, a world of paths, proximity and wayfaring. Spatial metaphors are one of the most widespread conventions in interactive system design, and are also carried across to the design of collaborative systems. However, in a collaborative setting, we are more directly faced with a need to pay attention to the social phenomena of space, and the social construction of meaning in spatial settings. The distinction at work here is that between "space" and "place", and the relationship between the two. Collaborative systems design has traditionally paid little attention to this relationship, although it has begun to receive some more attention lately.

In this section, we shall consider this relationship, starting from the position, again, that social navigation is a phenomenon of interaction. This section will briefly recap on its consequences for the design of collaborative systems, and then consider how social navigation fits into the picture when the notion of space has been reformulated.

3.1 On "Space" and "Place"

Collaborative systems are frequently organised around computational spaces of some form. We encounter systems employing media spaces [8], shared workspaces [34], argumentation spaces [40], etc. The idea of "space" is widespread because it is so fundamental to our everyday experience. Our conception of the world is fundamentally spatial; our own three dimensional embodiments in the world are the most fundamental part of our everyday experience. The use of space as an organising metaphor for interaction (and, indeed, for many other things besides [35]) is a natural one.

On the other hand, and just as we found earlier in considering navigation in shared virtual reality systems, the idea of "space" is one that deserves some consideration. Spatial settings are conveniently familiar and all-encompassing, but just what role is spatiality playing when it is adopted as an interactional metaphor? Further, what role does it play when the interaction is collaborative or social in nature?

Interface developers have only been in the space business for a couple of decades; perhaps it's not surprising that our view is relatively unsophisticated. So, in trying to look at these questions, one place to turn is to the Built Environment (architecture, urban design, etc.), where issues of space, of interaction and of design have been combined for thousands of years. Drawing on architectural and social theorising, I have argued, elsewhere, for a reconsideration of the notion of "space", and in particular, a reassessment of the relationship of "space" to "place" [25]. The argument presented in that paper was rooted in the philosophy behind the design of "media space"

technologies. We observed, in particular, that media spaces had been designed around an understanding of the relationship between the structure of the environment and emergent understandings of the action that takes place there. It is no coincidence that the original developers of media space technologies had backgrounds in architecture. The technology was dubbed media "space" precisely because its design took an emergent view of the relationship between space and place.

What is this relationship? At its most primitive, it's the relationship between structural and social aspects of the designed environment. *Space* refers to the three-dimensional structure of the world, and the configurations of light, air and material that create lots, buildings, rooms, conference centres, churches, theatres, casinos, shopping malls, offices, nooks, parks, and the various other familiar configurations of spatial setting with which we are familiar in the everyday works. Alongside this world of spatial settings is a world of places. While spaces take their sense from configurations of brick, mortar, wood and glass, places take their sense from configurations of social actions. Places provide what we call *appropriate behavioural framing*; on the basis of patterns of social action and accountability, places engender a set of patterned social responses. Spaces provide physical constraints and affordances, based on things like the fact that it's easier to go downhill than up, that humans cannot walk through walls, and that light passes through glass. In parallel, places provide social constraints and affordances, based on things like the fact that Western society frowns on public nudity, courts and churches are places for more dignified affairs than nightclubs, and that joyful exuberance is an acceptable response to sporting events but not to conference presentations.

Until recently, this relationship has, largely, gone unexamined in the CSCW literature. On the other hand, the emergence of collaborative virtual reality systems and observations of the natural and compelling nature of interaction in spatial settings. This has resulted in a variety of recent explorations of collaborative and interaction systems from the perspective of architectural and urban design, including the work of Erickson [18], Fitzpatrick et al. [20], and Benyon and Hook [7], as well as the contributions of Chalmers and Dieberger in this volume.

Space and place are fundamentally intertwined, of course; they each influence and condition the other. However, they are distinguishable. The design of the built environment is precisely about the artful manipulation of the relationship between space and place, between physical structure and social action (and, of course, about the history of the relationship between the two).

In his book "City: Rediscovering the Center" [44], William Whyte presents a detailed study of the everyday elements of urban life. His concern is with the functioning of urban spaces, in particular the densely populated daytime downtown areas. On the basis of photographic and statistical studies, he builds up an image of spaces that "work" and spaces that "don't"; ones which succeed or fail at creating a sense of place. He documents where people like to stop and have conversations, or eat their lunch; which spots attract crowds and which are deserted. Throughout his discussion of patterns of activity on streets, at plazas and in parks, he emphasises the idea that "what attracts people is people." The role of space is to frame human action.

3.1.1 Space and Place in Collaboration

What can we learn from this? What would such a reconsideration say about collaborative systems?

The functioning of virtual spaces and physical spaces are clearly very different (it's all too easy to forget that the relationship is metaphorical), but we can find evidence for similar relationships between space and place in a virtual setting. Harrison and Dourish present a number of examples from the research literature of virtual spaces working or not working. For instance, experiments in the use of video technology to link public spaces for the purposes of informal communication

have had very mixed results. While some, such as the PARC/Portland link reported on by Bly et al. [8] have been extremely successful and popular with participants, others such as the Virtual Window discussed by Fish et al. [19] have languished largely unused.

We can speculate about the reasons, and the potential consequences for design; clearly, the set of issues surrounding the deployment of these sorts of technologies are extremely complex. On the other hand, one clear indication is that the immediate social context is important. The issue is how the virtual space that the technology creates comes to be peopled and inhabited, how people come to have an understanding of what it does for them and how they should act. Harrison and Dourish point particularly towards a notion of *appropriation* – the extent to which the technology lends itself to be taken over by the participants and turned to their own uses, so that they can structure the space around their own needs and activities and make the technology their own. In other words, what we find to be important is not the technology itself, not the "space" that it creates, but the creative peopling of that space to turn it into a place, where people do things.

The other significant consequence for the development of collaborative systems echoes an earlier observation. It was noted earlier on in contrasting social navigation with spatial and semantic forms that social navigation can take place both in and out of spatial settings. It may be that social navigation is effected in a spatially organised environment, but it may also happen in non-spatial settings. Analogously, at a conceptual level we may find "places" that are non-spatial at heart (or, at least, ones whose spatial nature has nothing to do with their functioning as places). One simple example is an electronic discussion group such as a mailing list, Web forum or USENET news group. These have many features that we associate with "place"; they have a set of behavioural norms and expectations that frame the activity of individuals and against which activity can be considered. However, they do not offer what we would think of as "spatial" features; there is no up and down, no near and far, none of the "technology" of the everyday world. They are space-less places. Space might be a convenient, compelling and familiar metaphor for the development of collaborative systems, but it is not fundamental, nor necessary to their social functioning.

This reflects the fact that what places (rather than spaces) offer is what we called appropriate behavioural framing. Dieberger [this volume] presents interesting evidence for the ways in which people have mutually held understandings of what he calls the "social connotations" of places; social connotations are the consequences of behavioural framing, and they operate in both real world and virtual environments. Discussions of appropriate postings to newsgroups or other forms of discussion forum reflect the same concerns in non-spatial settings.

So these two concerns – appropriation and appropriate behavioural framing – help to set up an initial framework for considering how the ideas of space and place apply to collaborative technologies in general. The next question to ask is how they help explain a relationship between space, place and social navigation.

3.1.2 Social Navigation from the Perspective of Spatial Settings

When we think about it from the perspective of spatial settings, we immediately encounter some interesting features of the term "social navigation" itself. Indeed, we are immediately confronted with two diametrically opposed positions, from which "social navigation" is either an oxymoron or a tautology.

Social Navigation as an Oxymoron. The first position is that "social navigation" is an oxymoron. After all, navigation is a phenomenon fundamentally rooted in the physical world, an arrangement of, and orientation to, the elements of the physical world. From this perspective, navigation is a purely physical phenomenon, and not a social one.

There would seem to be at least two responses to this. The first is that, although navigation in the everyday world might be a purely physical phenomenon, there is no physical instantiation of the

virtual information spaces we are concerned with when applying "social navigation" in electronic settings, and so the support of social features is more pressing requirement in virtual environments. The second is that the purely physical construal of everyday navigation is a flawed one, and that, in everyday settings and the physical world, our navigation and wayfaring are actually highly socially-conditioned phenomena. This response leads us to the second position.

Social Navigation as a Tautology. The second position is that "social navigation" is a tautology. From this position, all navigation behaviour is social action and fundamentally conditioned by social forces. We can follow up this observation in a number of different ways.

One example is to take "navigation" itself, as a practical matter. Hutchins has developed this position in considerable depth in a number of investigations [32, 33]. His particular concern has been the practice of navigation on board large ships. He has produced a detailed analysis of the practice of ship-board navigation, from the perspective of distributed cognition, observing how the measurements and calculations which comprise navigation on ships are distributed over a group of individuals working together. Navigation is not so much performed as achieved through a complex of interlocking individual activities.

The tools of navigation are, themselves, socially constructed. Maps, for example, which might seem to be objective representations, faithfully detailing the structure of the environment, are rife with social and political concerns. Who draws the maps? What ends do they serve? Who decides what constitutes a "permanent feature", what is or not named, what is or is not mapped and recorded? What is the reality of the straight black lines that delimit one region from another? Wood [45] presents a wide-ranging analysis of the power of maps, arguing that maps construct the world rather than representing it.

The other tool of navigation, of course, is the world itself, but even that is not straight-forward. Lynch [36], for example, discusses the "imageability" of cities. Not all places are equally navigable, and the actual practice of navigation depends on specific features of the environment. Chalmers et al. [10] has explored how imageability features could be incorporated into virtual information spaces. However, what those features might be, what constitutes navigability and how navigation takes place, none of these are human constants. A number of social scientists, including Suchman [41] and Hutchins [31] have explored the metaphor of the differences between Western and Polynesian navigational practices, which have allowed Polynesian sailors to navigate, without instruments, long voyages inconceivable to western observers.

3.2 Place for Social Navigation

The distinction between "place" and "space" provides a frame for analysing the structure of, and action within, computational environment for collaboration. How does this frame relate to social navigation? There are at least two ways we can see a relationship between these ideas. The first is the question of how a space becomes populated; and the second is how social activity might transform it.

If it is sense of "appropriate behavioral framing" that distinguishes, conceptually, between space and place, then the most important factor is the presence and activity of people. In other words, the fundamental benefit we gain from social navigation is that it is a way of populating the information space. Social navigation, after all, hinges on two fundamental features; first, the presence of multiple individuals within some space, and second, the communication of aspects of their activity to each other (which we called "awareness" in the earlier discussion).

By seeing something of the activity of others, I can gain an understanding of the behaviour of individuals in a space, and so gain a sense of the style of appropriate action. When the space is populated, it becomes invested with a sense of appropriateness. That sense of appropriateness is, of course, a phenomenon which emerges from the activity of the individuals themselves; it is

subject to change over time, and continually evolves around the contents of the space, the proclivities of the people there, the affordances it offers, and so on. It becomes a place, with a set of understood behaviours, norms and expected practices.

The power of social navigation, or the foundations upon which social navigation is built, is to give the space meaning. That meaning comes from the collective sense of the action that takes place within the space, which itself is captured, made manifest and communicated by social navigation technologies. Since a user can see something of what has happened in the space before, then they can gain a sense of the history of the space and hence gain a sense of a set of spatially oriented practices. The book recommendation system at an on-line bookstore provides gives someone a sense of what other texts are read by people who are interested in a topic that they are investigating, which in turns gives them a sense of which authors are prominent and respected in the area (and by extension which ones are not), which books are regarded as definitive, and even, perhaps, of what other people find this topic interesting, other things they like, and so forth. The space of books has become populated by other people, and an individual's activities within the space are guided by a sense of how that space is currently configured and inhabited. Recently, others such as Benford et al. [6] have begin to investigate the opportunities for considering the World Wide Web as a space to be populated, so that the space itself begins to reflect something of the actions that take place within it.

The other side of the coin is the impact of the user's action on the space. In other words, if a place is a setting for action, how is the setting influenced or changed by the action that takes place there? One form of this impact is, again, the domain of workspace awareness technologies described earlier; for example, the "edit wear" approach is a way of leaving a mark in the space as a sort of "computational erosion".

One particularly interesting place to look at the way in which a place can not only frame action but also serve to represent and communicate it is in MUDs, persistent text-based virtual reality environments [13]. Most modern MUDs are programmable, allowing users to create new objects and behaviours that can be left in the space, cloned, copied, given to other users, used generally, and so forth. One facet of this is that a common way of learning the programming language that supports the development of new MUD objects is to examine and change other objects. In this way, programmed objects become units of exchange in what MacLean et al. [37], in a different context, have called a "tailoring culture". The MUD provides the mechanisms that support the development of this culture by allowing customisations to be made explicit and shared. What is particularly interesting about this setting, however, is that the MUD is also a place of social interaction and behaviour. This has two effects of the development of programmable artifacts in MUDs. The first is that the "objects" that are created are not simply artifacts, but can also be new modes of behaviour (new capabilities for players achieved through the creation of new "verbs" in the available lexicon) and new responsive places for interaction to happen (new "rooms" with special abilities). The second is that the patterns of social interaction in the environment set a context against which the development of these new programmable artifacts is set. For example, Cherny [11] presents an analysis of conversational behaviour in a MUD and documents the emergence of specific conversational patterns that are part of the established culture of that specific environment. One interesting feature is the way that, although these patterns of conversational behaviour emerge independently, they consequently became embodied in the user-configurable *technology* of the environment, and, hence, became available and inspectable to the participants in that environment. In other words, the place itself changed to reflect the commonly occurring patterns of activity within that space, and in such a way that aspects of those behaviours could be "read off" the space.

This ability to reflect patterns of activity such that they can be "read off" the setting in which they occur is, clearly, a generalised form of the approach taken by social navigation systems. The

perspective of “space” and “place” gives us a means to understand the role that the setting, and in particular, the setting as a populated site of social interaction, plays in making these connections.

4 Conclusions: New Opportunities for Social Navigation

The concept of “social navigation” has come of age. Since early observations about the role that could be played in information systems by information about the activities of others, and how, through that channel, individual use of information systems could be enriched more directly with social practice, these forms of systems have become familiar, accepted and even commonplace. In particular, with the rapid spread of the Internet and networked information systems, the use of social navigation, principally but not exclusively in the form of recommendation systems, has become integral to how we consider asynchronous interpersonal interaction and the management of socially-organised behaviour in virtual information environments.

It is time, then, to step back and consider just what social navigation *is*, how it works, and what it means. It is important to recognise that social navigation is not a sort of technology, but rather is a phenonemon of interaction. In particular, it is fruitful to consider how it relates to other perspectives on collaborative computational practice. In this chapter, I have been concerned, firstly, with these sorts of reflections and, secondly, with what we can learn from them for the design of new technologies.

This chapter has focussed, in particular, on two current perspectives in collaborative systems. One is the perspective of collaborative awareness. The notion of collaborative awareness as a fundamental concern in the development of (especially synchronous) CSCW tools has a longer history than social navigation, but shares a great deal in common with it. What it particularly interesting is how the two areas of research have taken different perspectives on the same fundamental issue of providing a means for individuals to discover and exploit information about others. The techniques which have been developed in social navigation systems, in addition, hold promise for tackling some of the problems with current approaches to awareness. We might hope, in the end, to see collaborative awareness and social navigation seen as aspects of the same phenomenon, and a unified design approach emerge.

The second perspective that has been explored here is that of the conceptual roles played by notions of “space” and “place” in collaborative settings, and how each of these influences the behaviour of individuals and groups in collaborative settings. One interesting feature here is that social navigation is built upon the same foundations that motivate a “place”-centric perspective on collaborative systems, one oriented around “peopled” spaces and a sense of “appropriate behavioural framing” that emerges from not the visibility of social conduct within a space. As such, then, the lessons learned from investigations of social navigation have an important role to play in explorations and further development of this place-centric view.

Along the way, various underlying concepts have emerged as features of the landscape. Social navigation systems support *aggregation*, *transformation* and *decoupling* of information, allowing them to present awareness information in terms of trends rather than specific actions. Populated social places offer *appropriation* and *appropriate behavioural framing*, distinguishing them from simple spaces, which are characterised in terms of their dimensionality.

What emerges from this is a picture of a reflexively populated information space. By “populated”, I mean that it contains not just information, but also people who are acting on that information, and who can see the effects of each others actions and exploit that information in managing their own activities. By “reflexively” populated, I mean that not only does the structure of the space have an impact on the action of the users, but that the action of the users can also have an impact on the space. The space is malleable, adjustable to reflect patterns of action and the needs of the users who inhabit it. It supports the forms of appropriation encountered in the discussions of

media space design and similar environments. The lesson of those experiences is that simply populating a space is not sufficient, but that appropriation and malleability are equally important. Similarly, the lesson of social navigation is that the social element of information seeking are critical resources in the development of collaborative systems, whether those are based on physical real-world metaphors or chart new virtual spaces. The explorations in this chapter have suggested that these lessons have some bearing on each other, and that social navigation has a role to play in the broader design of interactive and collaborative systems. At the same time, considering the role that social navigation can play in other sorts of collaborative setting has prompted reflections on underlying principles that may, in the future, help to frame a theoretical account of the mechanisms supporting socially-supported information seeking.

Acknowledgements

Consideration of these issues over the years has been hugely enriched by stimulating discussions with a variety of colleagues, particularly Victoria Bellotti, Matthew Chalmers, Tom Erickson, Bill Gaver, Steve Harrison and Christian Heath.

References

1. Ackerman, M. and Starr, B. (1995). Social Activity Indicators: Interface Components for CSCW Systems. Proceedings of the ACM Symposium on User Interface Software and Technology UIST'95, 159-168. New York: ACM.
2. Baecker, R., Nastos, D., Posner, I., and Mawby. K. (1993). The User-Centered Iterative Design of Collaborative Writing Systems. Proc. INTERCHI'93, 399-405. New York: ACM.
3. Bellotti, V. and Sellen, A. (1993). Design for Privacy in Ubiquitous Computing Environments. Proceedings of the European Conf. Computer-Supported Cooperative Work ECSCW'93. Dordrecht: Kluwer.
4. Benford, S. and Fahlen, L. (1993). A Spatial Model of Interaction for Large Virtual Environments. Proceedings of the European Conf. Computer-Supported Cooperative Work ECSCW'93. Dordrecht: Kluwer.
5. Benford, S., Bowers, J., Fahlen, L, Greenhalgh, C., and Snowdon, D. (1995). User Embodiment in Collaborative Virtual Environments. Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'95, 242-249, New York: ACM.
6. Benford, S., Snowdon, D., Brown, C., Reynard, G. and Ingram, R. (1997). The Populated Web: Browsing, Searching and Inhabiting the WWW Using Collaborative Virtual Environments. Proceedings of INTERACT'97, 539-546. Chapman and Hall.
7. Benyon, D. and Hook, K. (1997). Navigating in Information Spaces: Supporting the Individual. Proceedings of INTERACT'97, 39-46. Chapman and Hall.
8. Bly, S., Harrison, S., and Irwin, S. (1993). Media Spaces: Bringing People Together in a Video, Audio and Computing Environment. Communications of the ACM, 36(1), 28-47.
9. Carlsson, C. and Hagsand, O. (1993). DIVE: A Platform for Multi-User Virtual Environments. Computers and Graphics, 17(6), 663-669.
10. Chalmers, M., Ingram, R., and Pfranger, C. (1996). Adding Imageability Features to Information Displays. Proceedings of the ACM Symp. User Interface Software and Technology UIST'96, 33-39, New York: ACM.
11. Cherny, L. (1995). The MUD Register: Conversational Modes of Action in a Text-Based Virtual Reality. Unpublished PhD dissertation, Palo Alto: Stanford University.
12. Donath, J. (1995). Visual Who: Animating the Affinities and Activities of an Electronic Community. Proceedings of the ACM International Multimedia Conference, 99-107, New York: ACM.

13. Dourish, P. (Editor). (1998). Interaction and Collaboration in MUDs. Special Issue of Computer Supported Cooperative Work, 7(1-2).

14. Dourish, P. and Bellotti, V. (1992). Awareness and Coordination in Shared Workspaces. Proceedings of the ACM Conf. Computer Supported Cooperative Work CSCW'92, 107-114, New York: ACM.

15. Dourish, P. and Bly, S. (1992). Portholes: Supporting Awareness in a Distributed Work Group. Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'92, 541-547, New York: ACM.

16. Dourish, P. and Chalmers, M. (1994). Running Out of Space: Models of Information Navigation. Short paper presented at HCI'94 (Glasgow, Scotland).

17. Engelbart, D. and English, W. (1968). A Research Center for Augmenting Human Intellect. Proceedings of the Fall Joint Computer Conference (San Francisco, CA). 393-410. Reston, VA: AFIPS.

18. Erickson, T. (1993). From Interface to Interplace: The Spatial Environment and a Medium for Interaction. Proceedings of COSIT'93 (Elba), 391-405.

19. Fish, R., Kraut, R. and Chalfonte, B. (1990). The VideoWindow System in Informal Communication. Proceedings of the ACM Conf. Computer-Supported Cooperative Work CSCW'90, 1-11, New York: ACM.

20. Fitzpatrick, G., Kaplan, S. and Mansfield, T. (1996). Physical Spaces, Virtual Places and Social Worlds: A Study of Work in the Virtual. Proceedings of the ACM Conf. Computer-Supported Cooperative Work CSCW'96 (Boston, MA), 334-343. New York: ACM.

21. Fuchs, L., Pankoke-Babatz, U. and Prinz, W. (1995). Supporting Cooperative Awareness with Local Event Mechanisms: The GroupDesk System. Proceedings of the European Conf. Computer Supported Cooperative Work ECSCW'95, 247-262, Dordrecht: Kluwer.

22. Goldberg, D, Nichols, D., Oki, B., and Terry, D. (1992). Using Collaborative Filtering to Weave an Information Tapestry. Communications of the ACM, 35(12), 61-70.

23. Greenhalgh, C. and Benford, S. (1995). Virtual Reality Tele-Conferencing: Implementation and Experience. Proceedings of the European Conf. Computer-Supported Cooperative Work ECSCW'95, 165-180, Dordrecht: Kluwer.

24. Gutwin, C., Greenberg, S., and Roseman, M. (1996). Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets and Evaluation. In Sasse, Cunningham and Winder (eds), People and Computers XI: Proceedings of HCI'96, 281-298, London: Springer.

25. Harrison, S. and Dourish, P. (1996). Re-Place-ing Space: The Roles of Space and Place in Collaborative Systems. Proceedings of the ACM Conf. Computer-Supported Cooperative Work CSCW'96, 67-76, New York: ACM.

26. Heath, C. and Luff, P. (1992). Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. Computer Supported Cooperative Work, 1(1-2), 69-94.

27. Heath, C., Jirotka, M., Luff, P., and Hindmarsh, J. (1994). Unpacking Collaboration: The Interactional Organisation of Trading in a City Dealing Room. Computer Supported Cooperative Work, 3, 147-165.

28. Hill, W., Hollan, J., Wroblewski, D., and McCandless, T. (1992). Edit Wear and Read Wear. Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'92, 3-9, New York: ACM.

29. Hill, W., Stead, L., Rosenstein, M., and Furnas, G. (1995). Recommending and Evaluating Choices in a Virtual Community of Use. Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'95, 194-201, New York: ACM.

30. Hindmarsh, J., Fraser, M., Heath, C., Benford, S. and Greenhalgh, C. (1998). Fragmented Interaction: Establishing Mutual Orientation in Virtual Environments. Proceedings of the ACM Conf. Computer-Supported Cooperative Work CSCW'98 (Seattle, WA), 217-226. New York: ACM.

31. Hutchins, E. (1983). Understand Micronesian Navigation. In Gentner and Stevens (eds), Mental Models. Hillsdale, NJ: Erlbaum.

32. Hutchins, E. (1990). The Technology of Team Navigation. In Galagher, Kraut and Egido (eds), Intellectual Teamwork: The Social and Technological Foundations of Cooperative Work, 191-221, Hillsdale, NJ: Erlbaum.

33. Hutchins, E. (1995). Cognition in the Wild. Cambridge: MIT Press.

34. Ishii, H. (1990). TeamWorkStation: Towards a Seamless Shared Workspace. Proceedings of the ACM Conf. Computer Supported Cooperative Work CSCW'90, 13-26, New York: ACM.

35. Lakoff, G. and Johnson, M. (1980). Metaphors We Live By. Chicago: University of Chicago Press.

36. Lynch, K. (1960). The Image of the City. Cambridge: MIT Press.

37. MacLean, A., Carter, K., Moran, T. and Lovstrand, L. (1990). User-Tailorable Systems: Pressing the Issues with Buttons. Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'90. New York: ACM.

38. Palfreyman, K. and Rodden, T. (1996). A Protocol for User Awareness on the World Wide Web. Proceedings of the ACM Conf. Computer Supported Cooperative Work CSCW'96, 130-139, New York: ACM.

39. Shardanand, U. and Maes, P. (1995). Social Information Filtering: Algorithms for Automating "Word of Mouth". *Proceedings of the ACM Conf. Human Factors in Computing Systems CHI'95*, 210-217, New York: ACM.

40. Streitz, N., Hannemann, J., and Thuring, M. (1989). From Ideas and Arguments to Hyperdocuments: Travelling through Activity Spaces. *Proceedings of the ACM Conf. On Hypertext, 343-364*, New York: ACM.

41. Suchman, L. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge University Press.

42. Tang, J., Isaacs, E., and Rua, M. (1994). Supporting Distributed Groups with a Montage of Lightweight Connections. *Proceedings of the ACM Conf. Computer Supported Cooperative Work CSCW'94*, 23-34, New York: ACM.

43. Terry, D. (1993). A Tour through Tapestry. Proceedings of the ACM Conf. Organisational Computing Systems COOCS'93, 21-30, New York: ACM.

44. Whyte, W. (1988). City: Rediscovering the Center, New York: Doubleday.

45. Wood, D. (1992). The Power of Maps, New York: The Guilford Press.

4.1.1 Biographical Details

Paul Dourish is a Member of Research Staff in the Computer Science Laboratory of the Xerox Palo Alto Research Center in Palo Alto, California. His primary research interests lie in the area of collaborative and interactive systems, and in particular in the patterns of mutual influence at work between social behaviour and system architecture and design. He holds a B.Sc. in Artificial Intelligence and Computer Science from the University of Edinburgh, and a Ph.D. in Computer Science from University College, London.

4.1.2 Contact Details

Paul Dourish, Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA
Email: dourish@parc.xerox.com
URL: http://www.parc.xerox.com/csl/members/dourish