Awareness Space in Distributed Social Networks

L.Gavin, S.Keuppers, C.Mottram & A.Penn
Bartlett School of Graduate Studies, University College London, UK

Key words: Algorithmic Form Generation, Distributed Workgroups, Space Syntax

Abstract: In the real work environment we are constantly aware of the presence and activity of others. We know when people are away from their desks, whether they are doing concentrated work, or whether they are available for interaction. We use this peripheral awareness of others to guide our interactions and social behaviour. However, when teams of workers are spatially separated we lose 'awareness' information and this severely inhibits interaction and information flow. The Theatre of Work (TOWER) aims to develop a virtual space to help create a sense of social awareness and presence to support distributed working. Presence, status and activity of other people are made visible in the theatre of work and allow one to build peripheral awareness of the current activity patterns of those who we do not share space with in reality. TOWER is developing a construction set to augment the workplace with synchronous as well as asynchronous awareness. Current, synchronous activity patterns and statuses are played out in a 3D virtual space through the use of symbolic acting. The environment itself however is automatically constructed on the basis of the organisation's information resources and is in effect an information space. Location of the symbolic actor in the environment can therefore represent the focus of that person's current activity. The environment itself evolves to reflect historic patterns of information use and exchange, and becomes an asynchronous representation of the past history of the organisation. A module that records specific episodes from the synchronous event cycle as a Docudrama forms an asynchronous information resource to give a history of team work and decision taking. The TOWER environment is displayed using a number of screen based and ambient display devices. Current status and activity events are supplied to the system using a range of sensors both in the real environment and in the information systems. The methodology has been established as a two-stage process. The 3D spatial environment will be automatically constructed or generated from some aspect of the pre-existing organisational structure or its information resources or usage patterns. The methodology must be extended to provide means for that structure to grow and evolve in the light of patterns of actual user behaviour in

the TOWER space. We have developed a generative algorithm that uses a cell aggregation process to transcribe the information space into a 3d space. In stage 2 that space was analysed using space syntax methods (Hillier & Hanson, 1984; Hillier 1996) to allow the properties of permeability and intelligibility to be measured, and then these fed back into the generative algorithm. Finally, these same measures have been used to evaluate the spatialised behaviour that users of the TOWER space show, and will used to feed this back into the evolution of the space. The stage of transcription from information structure to 3d space through a generative algorithm is critical since it is this stage that allows neighbourhood relations to be created that are not present in the original information structure. It is these relations that could be expected to help increase social density.

1. PROJECT FRAMEWORK

In a co-located team, members typically learn from a wide range of cues about the activities of the other members, about the progress in the common task and about subtle changes in group structures and the organisation of the shared task environment. Most of this group awareness is achieved without overhead effort. A distributed (virtual) team – even if its co-operation is based on state-of-the-art groupware systems – today has a far lower level of awareness and opportunity for spontaneous, informal communication. This reduces the effectiveness of the joint effort, and makes co-operation a less satisfying experience for the team members. The TOWER project aims to bring the wealth of clues and information that create awareness and cohesion in co-located teams to the world of virtual teams and to present them in a Theatre of Work.

TOWER supports group awareness and chance encounters through a 3D environment that is at the heart of the Theatre of Work. Avatars and their symbolic actions represent users and their current actions on shared objects while using a groupware application. Avatars of users who work in a similar context appear spatially close in the 3D environment. The integration of the Theatre of Work into a work setting opens new opportunities for spontaneous encounters and spontaneous contacts.

The project aims to create an awareness infrastructure that collects and manages sensor signals from the real work environments of the team members and events from the team's information space, and distributes them to trigger indicators appropriate for the team members;

1.1 Project

This paper will set out a methodology for the design of dynamic 3D worlds where architecture and objects represent the task-related information spaces of a team and their dynamics. Each work environment is essentially unique. Recognising this, we will develop a construction set approach to our system that recognises the signs of democratic communication ‘pathmaking' to build a best-fit solution. This construction set enables an automatic creation and dynamic recreation of the virtual space based on a transformation of the elements of shared information spaces into elements and places of the Theatre of Work. A Space Module will be developed based on advanced spatial analysis techniques and an understanding of conceptual and physical distance to allow this construction set to build appropriate work environments.

The methodology has been established as a two-stage process. The 3D spatial environment will be automatically constructed or generated from some aspect of the pre-existing organisational structure or its information resources or usage patterns. The methodology must be extended to provide means for that structure to grow and evolve in the light of patterns of actual user behaviour in the TOWER space.

We have developed a generative algorithm that uses a cell aggregation process to transcribe the information space into a 3d space. In stage 2 that space was analysed using space syntax methods (Hillier & Hanson, 1984; Hillier 1996) to allow the properties of permeability and intelligibility to be measured, and then these fed back into the generative algorithm. Finally, these same measures have been used to evaluate the spatialised behaviour that users of the TOWER space show, and will used to feed this back into the evolution of the space.

The stage of transcription from information structure to 3d space through a generative algorithm is critical since it is this stage that allows neighbourhood relations to be created that are not present in the original information structure. It is these relations that could be expected to help increase social density.

In order to apply Space Syntax theory to the virtual world and its inhabitants, some specific data is required mapping the users, documents and their interactions. The methodology currently applied in 'real world' situations examines building layouts and maps the movement data, with this information the theory can be used to evaluating different strategies for changing the layouts in order to maximise any chosen interactions.

The principal data required for Space Syntax analysis is a comprehensive geometrical description of the environment and a log of the movements and interactions within it. In the application of Space Syntax to the Tower environment, the situation is altered somewhat by the possibilities offered by virtual environments as regards changing, or morphing, geometries. The rules associated with movements and interactions can now be used to morph any elements (avatars, documents and geometries) within the environment in order to instigate any required interaction.

Each element has a series of attributes. All elements have attributes relating to their origin, destination and subject, some have 'length' (documents), time (avatars) etc. Each element, in an effort to represent itself, morphs its movements and geometries according to certain rules. The rules relating to visibility and adjacency determines an optimum position for the element within the world at any given time. The aim is to have the elements organise and re-organise themselves in a pattern using rules based on information from the communication log. Thus the world adapts and changes to reflect the preoccupations of the inhabitants and forms an environment within which location can take on meaning.

2. THE SPACE MODULE

The intention within the TOWER space is to build loosely on the way that social interaction and generation take place in the real environment. Here we summarise the main operative factors we need to take into consideration. First, the fundamental intention of TOWER is not to replace the real environment, but to supplement it. In particular it is to provide a means of increasing social density for organisations that of necessity are spatially fragmented and so sparsely distributed in real space. There are two main reasons why social density is important to the work organisation:

First, density makes possible communication and information flow. Low-density organizations face considerable overheads in achieving adequate information flows.

Second, density directly affects the efficiency with which organisations are able to generate new structures and divisions of labour in response to changes in their operating environment. Low-density organisations will tend to fossilise and find it difficult to adapt to change.

Already the trend in organisations is to even greater spatial fragmentation made possible by the growth in communication technologies. Although these technologies make it possible for organisations to function whilst more spatially dispersed than ever before, the effect of dispersion is negative overall. TOWER aims to help redress the balance by developing a new channel for creating spatial awareness and social density.

2.1 Social Density

In real environments we know that a substantial proportion of communications take place in a random and relatively unpredictable way. People working in a concentrated fashion are considered to be unavailable even if they sit in an open plan space, but as soon as they get up from their desk to go to an urgent pre-planned meeting they are regarded as available by those that they pass and are recruited into conversation. Our studies suggest that up to 80% of work related conversations in the office take place in this unplanned and essentially unplannable way. It is unplannable since you cannot predict that any individual will walk past at any particular time (Penn et al 1997). Now, consider that most work processes require people to construct certain patterns of movement between people, information sources and spatial resources in the work environment. In an organisation with a division of labour into different tasks and processes this means that different people will have different characteristic movement patterns and space use regimes. A single building will structure these movement patterns and certain spaces will be used in common between different individuals (everyone comes through the entrance in the morning, or uses the lift and stairs regularly, or the coffee machine, photocopier or toilets). However, just as importantly, building layout also separates groups from each other. Hospital layouts, for example, are very effective in separating doctor's movement patterns from those of patients except in the formal spaces where the medical interface is planned to take place, and quite a degree of sophistication is needed in hospital planning to achieve this. In corporate offices often clients and top management are separated from workers.

In social terms, the combination of work processes and the regular movement paths that those entail, with the spatial layout of the building, create a pattern of co-presence between different categories of building user (and co-absence between other categories). These patterns of co-presence are fundamental to the reproduction of the social form of organisations, since, by and large, you need to know people and who they are before you phone them. It is thus the people that you know since you see them regularly, even if you have never worked with them directly, that become a main resource for generation of new links in organisations. We call this field of co-presence of people 'social density'. Of course most organisations have a number of other methods of achieving density, including brainstorming, management training weeks, retreats and so forth, but all these are about spatialising groups of people that management think need to be put in touch. What built space can do is put people in touch that management could not have predicted needed to talk. What built space can also do is to reproduce only the existing organisational structure and eliminate other links – this is a recipe for a fossilised organisation. The problem with dispersed organisations is that they tend towards fossilisation since active investment must be made in communication to overcome distance, and quite naturally, that investment is made on the rational basis of making existing links stronger, rather than supporting new links. This is where TOWER's main contribution will lie. By spatialising the organisation (albeit in a virtual awareness space) it will help to increase social density.

2.2 How do we achieve this?

The initial idea that we have been working with is that we need a two-stage process. We need to automatically construct or generate a 3D spatial environment from some aspect of the existing organisational structure or its information resources or usage patterns. We then need to provide means for that structure to grow and evolve in the light of patterns of actual user behaviour in the TOWER space. However, there is a risk that if we construct a spatial environment based directly on the information structure of the existing organisation, it will reproduce co-absence rather than generate new patterns of co-presence. A direct mapping of say the file hierarchy of the shared database would be expected to do this as file hierarchies are essentially treelike and to get between neighbouring branches one needs to pass through the trunk. We think that we need an indirect mapping therefore from information space to virtual space, and this mapping needs to generate virtual space with some quite distinctive spatial properties:

- it should be permeable everywhere (and in the main have few dead ends).
- it should be intelligible – that is have good local to global correlations.
- it should ideally both spatialise local groups and create larger scale relations between these groups.

An example of the spatial relations described in the last point could be viewed in terms of the distinctiveness of urban areas such as Covent Garden, Soho and Mayfair in London. You know when you are in each but cannot quite say where the boundary lies. In urban space this is achieved through the intelligibility structure of local to global relations.

The structure of the Space Module is based on a generative algorithm that uses a cell aggregation process to transcribe an information space into a 3D space. In stage 2 that space is analysed using space syntax methods to allow the properties of permeability and intelligibility to be measured, (the algorithm is explained in greater detail in later chapters and is referred to as 'culling') and then these are fed back into the generative algorithm. Finally, these same measures are used to evaluate the spatialised behaviours that TOWER users demonstrate in the space, and these are fed back into the evolution of the space .

The stage of transcription from information structure to 3D space through a generative algorithm is critical. It is this stage that allows neighbourhood relations to be created that are not present in the original information structure. It is these relations that could be expected to help increase social density. The feedback stage is also critical, since it is this feedback that will allows those neighbourhood relations that turn out to be useful to flourish and will eventually eliminate or replace those that turn out not to be useful.

3. CELL AGGREGATION

The primary functions of the Space Module are:

– to handle the generation of the 3D environment given an existing information infrastructure, and
– to handle the evolution of that environment based on patterns of use in the TOWER environment.

Both of these functions depend on the input information. An Event Notification Structure has been established in order to capture, distribute, and notify user operations or object state in both real and virtual work environments. The model is based on sensors, events, and indicators. Both sensors and indicators are part of the work environment. Sensors report user operations or object state encoded as events to the ENI server. The server accepts, collects, and processes these events. Indicators are used to visualise events. There are two modes for indicators to get notified about events: in the first mode the indicator ask for events and in the second mode the indicator is notified of events by the server.

The description of the existing information infrastructure is derived from logs of actual usage of information sources. By tracking time-based patterns in that usage, inferences are drawn on likely relevance of different items of information. These are used to give shape to the spatial relations between information resources as represented in the TOWER world. In this sense the environment will grow according to actual usage patterns of the underlying information resources.

The main paradigms that are being used are aggregate models of environment construction. In these, as each information resource appears in the ENI event database a spatial representation is created within the TOWER world. In its simplest form a special agent navigates the entire information space and creates special ENI events at each node so that the event database comprises a full description of the whole information space. Simple sets of rules govern the location and form of the representation and the way that it is related spatially to other information resources. These rulesets implicitly define the structure of space left over in the environment between each of the information resources. The precise form of the aggregation ruleset is therefore critical to the definition not only of local relations between different resources, but also of the global spatial configuration of the whole environment. However, both local relations and global configuration are not directly coded into the ruleset, but emerge dynamically as a result of the application of the ruleset to a specific log of information usage. Given a different usage pattern, a differently configured environment would emerge.

There are two main functional requirements of the rulesets. First, it is important that the local relations between information resources be 'sensible'. This means that there should be some relatively obvious logic to the local disposition of resources with respect to each other. Second, the whole configuration of the resulting aggregation should be relatively intelligible. It is our expectation that by aggregating according to an actual history of usage the local relationships will emerge as an aspect of the process. Thus information resources that are used in temporal proximity in an actual log will be located in spatial proximity in the TOWER environment. Our definition of 'intelligibility' of a spatial configuration is based on the space syntax concept of the prediction of global position from local information. Aggregation processes are being developed that lead naturally to relatively intelligible configurations.

The simplest of these aggregation processes is the 'beady ring', settlement form identified by Hillier & Hanson (1984). In this aggregation a closed cell open space dyad forms the aggregation unit. The rule is that open space must attach to open space. The first dyad is located and oriented at random. The next can take one of three open faces to attach its open space, and each of those sites can be connected in three possible orientations. Again each of these is selected at random. The process continues until all information resources have been located. The problem with this form of aggregation is that although relatively small aggregations are syntactically intelligible, as aggregations grow the relationship between local and global spatial properties breaks down.

3.1 Culling

For this reason a further stage of 'information culling' has been included in the process. In this stage, the environment is reviewed in terms of the intervisibility of different information resources. Information resources are rated according to measures of similarity or relevance, and these culling procedures seek to maximise the intervisibility of those resources that are most strongly related. This process tends to lead to a relocation of certain resources within the environment, and to the removal of resources from locations where they block the intervisibility of other resource pairs.

4. RULESET

4.1 Identifier and Descriptor Attributes

Each object in the world has two sets of attributes. The Identifier Attributes are automatically set by ENI events. They can include any 'field' in the ENI event list e.g. 'creator', 'reader', 'length of document', keywords. The second set of attributes are the Descriptor Attributes. These could take the form of the settings for descriptors such as the geometric form of the representation (these are currently cubic 'building blocks' in the early prototype implementation, but could be cylindrical 'trees' or any other geometric representation), colour, height, transparency, hue, and text labels.

Decisions have to be made regarding the mapping of identifier to descriptor attributes. Currently the early prototype world reflects the frequency of people accessing a document (summed from 'creator', 'reader', 'modifier' and 'date' fields in the ENI event record) represented by the height of the building block. There are, however, a number of other event attributes that we might choose to represent, such as MIME type, document size, creator, modified date etc. There are also a number of representation factors other than height, including change of geometry, colour, transparency, textual labelling or a combination of these. Location can also be considered a representational factor, however this is treated separately in the system since it is qualitatively different in that it allows the representation of the relationship between different events or information resources. The Attribute Matching System handles location.

4.2 Attribute Matching System

The Attribute Matching System is used to compare existing objects with a new object in order to decide where the new object should be placed within the world. Currently the early prototype matching system allocates points purely on the basis of similarity of Identifier Attributes. This assessment is carried out for each new object against each existing object in a 'snap' (card matching) type game e.g. the new document was 'created by Smythe' and the existing document was 'created by Smythe' would gain the match 5 points. If both documents were 'read by Prinz' the match would gain another 5 points. Currently if there is a free space next to the existing document this gains an additional point. The new object is eventually placed in a position next to an existing object with the highest match points and a free space.

The points allocation structure in the Attributes Matching System is an area in which we anticipate further development taking place. Users may wish to prioritise certain attributes and or give others negative points. The main issue here is to develop methods for location of information resources or event representations in which locations become meaningful to the user. This will inevitably differ from one end user application to another, and so the approach adopted at this stage is to develop a generic method for determining location based on attribute matches, where the specific attribute values can be user defined.

4.3 Ideal Neighbours and Actual Neighbours

Ideal Neighbours are the objects that score the highest points through the Attribute Matching System. Actual Neighbours are the objects that score the highest points through the system and have a free space and so are chosen as neighbours. Ideally, the Ideal Neighbour will be the Actual neighbour, however as a TOWER world grows free spaces may get used up and less than ideal locations for new blocks will be used. These definitions are utilised in the 'Culling Routine' in order to re-arrange the world through best Ideal Neighbours and sight lines.

4.4 Placement Rules

Once the Actual Neighbour has been identified the placement rules are followed. Each object is currently represented in the early prototype as a 'building-block' of enclosed (private) space with an adjoining 'plot' of open (public) space. The placement rules require neighbouring open plots to meet as in Figure 1. The cumulative effect of these placement rules creates blocks of private spaces, and streets and squares of public spaces. The effect is to generate a continuously accessible pattern of open space, irrespective of the nature or sequence of events and attribute matches describing the information space. It is impossible to generate areas of inaccessible open space.

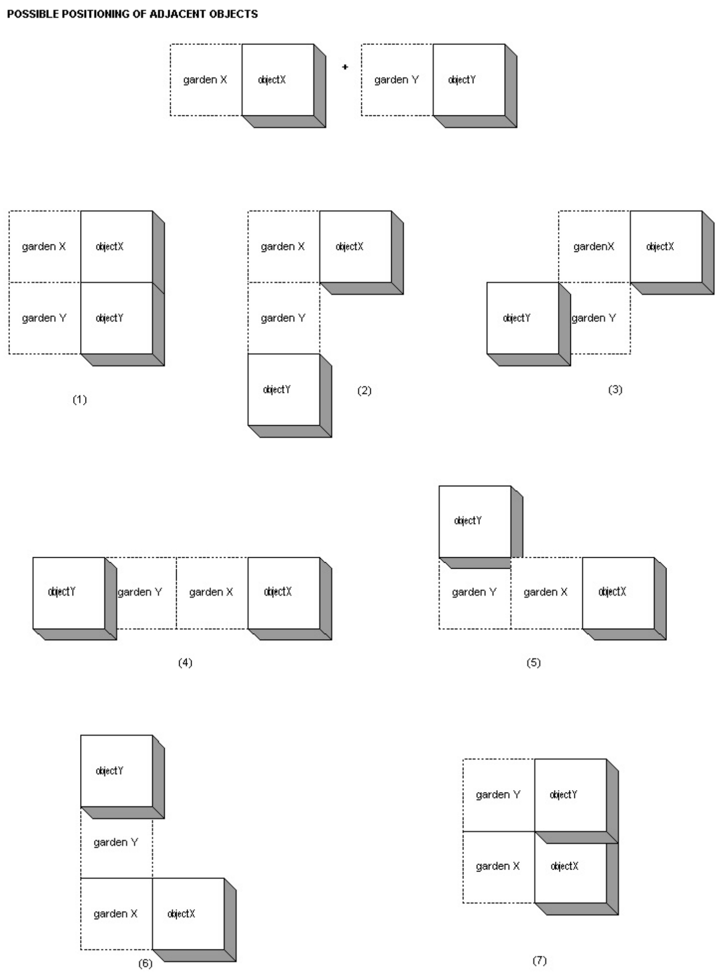

Figure 1. Placement Rules

The resulting patterns of development can be shown to have quite specific proto-urban qualities. All areas of open space relate directly to building blocks and as a consequence of building blocks related by attribute matching are likely to be visible to one another. As the space pattern grows above a relatively small size, local and global spatial properties are found to become correlated and this gives the spatial pattern a natural property of intelligibility. As the aggregation grows still further, however, intelligibility diminishes, and this calls for further processes to ensure that the world remains scalable. The main means of achieving this is the culling process.

This not only ensures that related items of information are inter-visible, but also serves to increase the degree of spatial interconnections between different parts of the evolving aggregation, and ensures increased intelligibility as a system grows larger.

4.5 Life Points

The Life points system helps to reflect a documents history within the world. Life points are added whenever a document is accessed, but subtracted whenever a 'culling' takes place. Presuming that the culling will take place on a regular basis, say daily, the document would eventually disappear if it was not accessed after a certain number of days. This process also allows for a degree of change in the layout over time.

5. RESULTS

5.1 Examples of the World Building and Culling Rulesets

Figure 2 shows the growth and culling process for a TOWER workspace using dummy event data. In each stage of the process new building blocks are added to the world representing new nodes in the database hosting the collaborative workspace. These are coloured according to context attribute values denoting degrees of similarity between items of information in the information space. The process of aggregation tries to bring information items that are closely related together in space, however as the world grows it becomes increasingly difficult to satisfy this requirement without rearranging the location of existing blocks. This is when the culling process comes into force. As time progresses the number of new information items reduces, and events increasingly involve hits on existing information items in the world. These events are represented in the prototype world as increases in the height of a block. Thus the most frequently visited information items become TOWERs in the landscape. As the world develops, the rules lead to a degree of structuring of the landscape, but not to a precise correspondence between layout and colour. This is a result of the fuzzyness of the context attributes and the fact that there are numerous competing associations between information items. The process leads to associations in space between items of information that were not adjacent in the original hierarchical file structure in the database. This is the key aspect of the TOWER world design that will lead to users being brought into contact with new and relevant items of information .

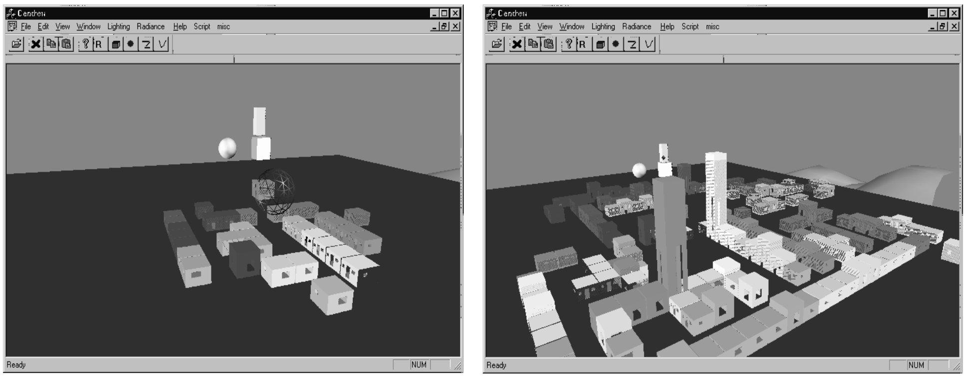

Figure 2. World Growth

5.2 The resulting spatial layout

The aggregation ruleset used to locate 'building blocks' in the TOWER world leads to a pattern of connections between open space in the world. This pattern is a characteristic of the rules used to aggregate and to relocate blocks during culling. Tests have been carried out using Space Syntax analysis techniques to evaluate the resulting spatial layout of the TOWER world. The intention is that as the information space grows it should maintain a degree of intelligibility. In space syntax analysis it is possible to provide a metric for the degree of intelligibility of a spatial system in terms of the degree of correlation between local and global measures of the graph describing intervisibility of points in space. Figure 3 shows measures of local (on the left) and global (on the right) measures for the TOWER world as it grows. Although there is a lack of correlation at the start of the process, as aggregation and culling proceed the result is that local and global measures become increasingly correlated. The final result of culling is to remove blocks in the centre in order to provide intervisibility between related information items.

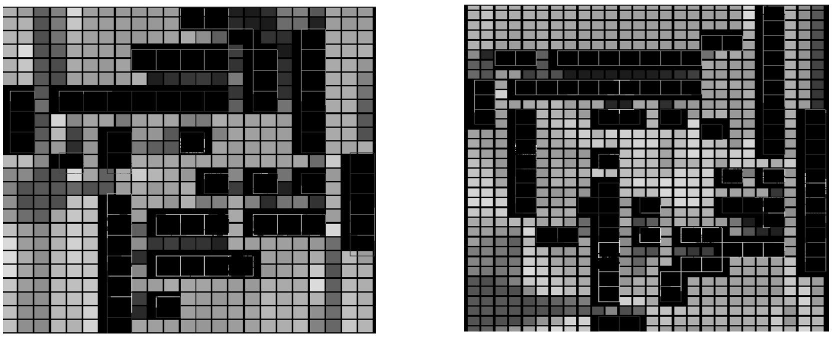

Figure 3. Connectivity Patterns

6. REFERENCES

Hillier, W. and J. Hanson, 1984, *The Social Logic of Space*, CUP, Cambridge.

Penn, A., J. Desyllas, J. and L. Vaughan, 1999, The Space of Innovation: Interactionand Communication in the Work Environment, in *Environment and Planning B:Planning and Design*, (1999), Vol. 26, 193-218, Pion, London.

Bates, Marcia J., 1996, "Document Familiarity in Relation to Relevance, Information Retrieval Theory, and Bradford's Law: The Getty Online Searching Project Report No. 5." in *Information Processing & Management* 32 (1996): 697-707.

Belisle C., R. Zeiliger, T. Cerrato, 1997, *Integrated Cognitive Engineering at the Interface: A Tool Mediation Perspective*, http://www.ish-lyon.cnrs.fr/labo/LIRE/CT97.htm.

Rose, A., W. Ding, G. Marchionini, J. Beale, and V. Nolet, 1997, *Building an Electronic Learning Community: From Design to Implementation*, http://www.cs.umd.edu/TRs/authors/Gary_Marchionini.html.

Thatcher, A., *Determining interests and motives in WWW navigation*, Psychology Department, University of the Witwatersrand, WITS, 2050, South Africa, http://cyberg.curtin.edu.au/members/papers/66.shtml