

Supporting Social Worlds with the Community Bar

Gregor McEwan and Saul Greenberg

University of Calgary

2500 University Drive NW

Calgary, AB, Canada T2M 1N4

+1 403 210 9501

{mcewan,saul}@cs.ucalgary.ca

ABSTRACT

The Community Bar is groupware supporting informal awareness and casual interaction for small social worlds: a group of people with a common purpose. Its conceptual design is primarily based on a comprehensive sociological theory called the Locales Framework, with extra details supplied by the Focus/Nimbus model of awareness. Design nuances are strongly influenced by observations and feedback supplied by a community who had been using both the Community Bar and its Notification Collage predecessor for a total of five years. As a consequence, Community Bar's design supports how communities of ad-hoc and long-standing groups are built and sustained within multiple *locales*: places that offer a group the site and means for maintaining awareness of one another and for rapidly moving into interaction. This includes a person's lightweight management of his or her membership in multiple locales, as well as ones varying engagement with the people and artefacts within them.

Categories and Subject Descriptors

K.4.3 [Computers and Society]: Organizational Impacts - *Computer-supported collaborative work*.

General Terms

Design, Human Factors, Theory.

Keywords

Casual interaction, Locales Framework, focus/nimbus, groupware.

1. INTRODUCTION

Various studies of white collar work sites report that a large portion of peoples' time is spent in unplanned, casual interactions with other collocated co-workers [10][16]. These interactions are stimulated by physical proximity: members of the group acquire informal awareness of each other, such as knowledge about presence, activity, and availability, and this knowledge leads to opportunities for people to engage in light-weight casual interactions at appropriate times and in an appropriate manner [10]. In contrast to formal meetings, casual interactions are unplanned, brief, frequent, and usually engage small groups of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GROUP'05, November 6-9, 2005, Sanibel Island, Florida, USA.

Copyright 2005 ACM 1-59593-223-2/05/0011...\$5.00.

people familiar with one another [16]. While seemingly mundane, these casual interactions prove important. They keep individuals informed about each other in social and professional contexts, they reinforce social bonds, and they make the transition to tightly-coupled collaboration easier [10][16]. These tightly-coupled collaborations easily take advantage of near-by work artefacts to progress naturally to artefact-centric work.

However, the same studies also found that these types of interactions are severely affected by physical separation, where there is an exponential drop-off in their number over even small distances such as that between offices at ends of the same hallway [10][16]. This means that distributed communities of co-workers miss out on these interaction opportunities. In response, groupware developers designed a myriad of informal awareness and casual interaction tools; each tool typically provides mechanisms for displaying informal awareness information that lead to casual interactions between distributed group members. Three popular examples are text-based Instant Messengers (IM) [12], chat rooms / MUDS [4], and video-based media spaces [1]. These tools, especially IM, have proved immensely valuable in practice. For example, while most IM systems provide only a rudimentary indication of other people's presence, even this minimal information is enough to create opportunities for textual chats. The lesson is that even minuscule awareness information combined with a crude communication medium is enough to trigger the casual interactions desired by a community.

Yet even the most widely accepted of these tools are shallow caricatures in terms of how they support the social practices of the individuals and groups that use them. Instant Messengers treat one's social communities as a disparate set of buddy lists, where they favour isolated chats between two people. Chat groups and their variants have rigid notions of how groups are defined, how one becomes members of it, how people present themselves to others, and how conversations are publicized. From a social science perspective, communities are far more dynamic than that.

Our goal is to create tools that go beyond this basic support of casual interaction. To achieve this goal, our design perspective is to ground development of casual interaction tools in social science theory. In particular, we are motivated by the Locales Framework [5], one of the few comprehensive theoretical group interaction frameworks in the computer science field, as well as the Focus and Nimbus model of awareness [14]. As we show shortly, we derive and combine principles from these theories and apply them to the design of the Community Bar, a groupware tool that supplies ad hoc groups with rich awareness information leading to casual interaction. We also show how this tool leverages and extends two previously introduced design ideas.

First, *media items* [9] are used as groupware building blocks to offer rich multimedia awareness and interaction capabilities. Second, these items are embedded within the *sidebar metaphor* [3], where a person sees awareness information of different groups at the screen's periphery, and can selectively drill down to more information and interaction. From this work, we will claim two contributions: firstly, the derivation of theory-based design principles for informal awareness and casual interaction systems, and secondly, the design and implementation of a system – the Community Bar – that follows these principles.

To set the scene, we first summarize the Locales Framework and the Focus and Nimbus model. Next we discuss design principles derived from these theories and from what is known about casual interaction. We then describe the design of the Community Bar tool and how it relates to these principles.

2. THEORETICAL BACKGROUND

Social Science is rich in theory, and many of these could be used to motivate groupware design. We decided to use the Locales Framework, for it is one of the few comprehensive sociological theories of group work specifically developed with groupware design in mind [5][11]. Yet we realized that this framework did not suffice by itself: because it is specified in a descriptive manner and at a very high level, we augmented its mutuality (awareness) aspect with Rodden's Focus and Nimbus model of awareness [14]. In this section, we briefly summarize both theories.

2.1 The Locales Framework

The Locales framework relies on two fundamental concepts. The first is the *social world* – a group of people with a common purpose. The common purpose can range from one that is formally defined, such as employees implementing a company's business model, to an informal purpose, such as friends who are maintaining their ongoing social relationships. The second fundamental concept used is *centres and peripheries*. The key principle is that members of the group have many different and shifting relationships; thus their involvement in the social world should be represented as a continuum rather than the binary 'on' or 'off' embedded in many groupware systems. For example, at an instant in time a group may be planning an event. Those very close to the centre may be involved in detailed organization. Another person, who may just attend the event, is somewhat more removed. Yet another may skip this particular event, so they are closer to the periphery of the group, at least for the moment. In real world situations such as these, boundaries are made only if required; in practice people can fluidly adjust their 'membership' from centre to periphery as a consequence of their interests and their actions. Of course, some social worlds have explicit rules, membership lists and duties that define people's roles and what they do, but these tend to be restricted to formal groups.

The framework is divided into five aspects that describe how social worlds behave [5], and that incorporate the concept of relationships as continuous rather than binary. While our summary below has a theoretical flavour to it, upon reflection the reader should realize that the framework articulates what regularly happens in everyday social life.

Locale Foundations. A *locale* is the site and means that a social world uses in its pursuit of the shared purpose. Sites are places

that the social world uses, and means are the objects within those places. An example of a locale for a collocated work group is their shared work room, where the site is the work room and the area around it, and the means include its furniture, whiteboards, markers, etc. Another locale they use could be an IM conversation; the site is the virtual conversation space and the means include the electronic artefacts, such as pictures, shared during the conversation. Social worlds typically use many different locales when engaging in their activities.

Civic Structure. No social world operates in isolation. Members are involved in multiple worlds at once. Social worlds exist within broader organisational structures, and sometimes smaller sub-worlds are contained within the social world. For example, the many social worlds within an organization (or social group) overlap and influence one another. Civic Structures describe the relevant outside influences on a social world.

Individual Views. As an individual engages in work, he/she is rarely involved in a single task to the exclusion of all others [5]. Individuals will simultaneously engage in multiple different tasks, across multiple social worlds. There are two important aspects to be considered; a *view* on one social world, and an individual's *viewset* across multiple social worlds. A view is how an individual sees a single social world (the people and the locales), and it is dependent on the level of engagement with the centre of that world. A viewset incorporates the individual's views of all the social worlds with which they are engaged, e.g., when juggling work and family tasks. The view across these tasks will change continually without fully switching out of any of them.

Interaction Trajectories describe the highly dynamic nature of social worlds. Social worlds have phases (e.g. setup, full operation, finalising). Awareness of past actions and outcomes, present situations, and visions for the future are important for creating plans and strategies. Interaction Trajectories describe how all five aspects of the Locales Framework changes over time. Locales will be set up, used so that the sites and artefacts are modified, and eventually discarded. Civic structures will change as relationships between social worlds evolve. Individual views and viewsets will constantly change as their focus changes and their relationships to others changes.

Mutuality. Awareness of people, spaces and resources is vital for collaboration within the social world [10][16]. Fitzpatrick teases apart the definition of mutuality into provision and reception of awareness information. Members of the social world make information about themselves and their activities available to others. Others 'become aware' as they perceive this information. The separation of information provision and perception is important, as not all provided information is necessarily perceived. The Focus and Nimbus model of awareness [14], described next, defines this aspect of mutuality in more detail.

2.2 Focus and Nimbus Model of Awareness

The Focus and Nimbus model explicitly breaks down awareness into an interaction between the observer and the observed [14]. Each person or artefact in the environment provides some perceivable information about itself, called the *nimbus*. Conversely, each person in the environment has capabilities to perceive this information, and the way in which they direct this perceiving capability is called *focus*. As Figure 1 illustrates, the

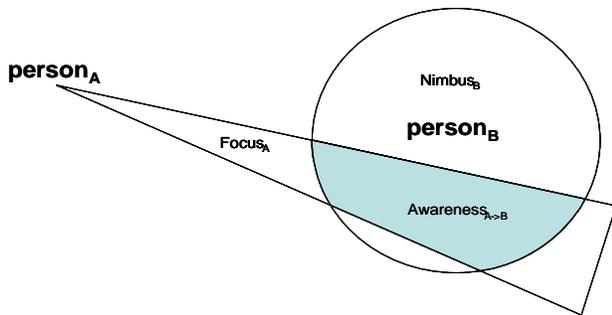


Figure 1: Focus and Nimbus combine to create Awareness

awareness that person_A has of person_B (or that person's objects or actions) is a function of the overlap of the focus of person_A with the nimbus of person_B. This model is fairly general. Applying it to analyse specific cases involves determining exactly how focus and nimbus interact to determine awareness. Consider people working around a table. Nimbus is represented by particular people's appearance, gestures, expression, voice, etc. One individual's focus will be concentrated on the people they are talking to, while less directed to others at the table. Of course, focus and nimbus is fluid, and changes over time.

The focus and nimbus model augments the mutuality axiom of the Locales framework in that: (1) awareness is defined by both the observer and the observed; and (2) awareness can be conceived as a continuous function rather than discrete.

3. DESIGN PRINCIPLES

The prior work on casual interaction and the above theories describe social phenomenon. However, they do not state criterion for groupware design. In this section, we transform the theory into groupware design principles. Lack of space precludes discussion of how each principle was derived from the theories, but the relationship between the two should be self evident. Most restate the major points discussed in the theories as design imperatives. Principles 1 to 4 are derived from the informal awareness and casual interaction literature, e.g., [10] [16]. Principles 5 to 8 are from Greenberg et al's transformation of the Locales Framework into heuristic evaluation principles for groupware [7]. Principles 9 to 11 are from the focus and nimbus model of awareness [14].

- 1 **Discretionary information should be visible at the periphery.** Awareness information needs to be constant and dynamic to maintain knowledge of the surrounding environment. However, it should not interfere with focus on other tasks.
- 2 **Allow lightweight transitions from awareness to interaction.** A primary benefit of having informal, peripheral awareness is as a basis for casual interaction. As casual interactions have to be lightweight, unplanned, and frequent, any tool that supports them must also reflect these properties.
- 3 **Support small groups of intimate collaborators.** Informal awareness and casual interaction are prevalent within small groups of people (~2-15) that know each other well [10][16].
- 4 **Provide rich information sources and communication channels.** Awareness can be based on many different cues. The more information that is presented, the better people are able to interpret awareness information. Casual interaction

also takes many forms and so there is a requirement for multiple ways of sharing and communicating.

- 5 **Provide centres (locales).** People work in multiple contexts simultaneously, switching between them. The design of informal awareness / interaction tools should provide centres or locales that comprise a site and means for the social world.
- 6 **Provide a way to organise and relate locales to one another (civic structures).** Locales relate to each other in different ways. A representation of an individual's locales needs to allow the relationships between the locales to be expressed.
- 7 **Allow individual views.** Each person interacts with a set of locales in different ways. As their level of engagement across their viewset changes, so they should be able to change the presentation content of locales as well as the objects and people contained in the locales.
- 8 **Allow people to manage and stay aware of their evolving interactions over time.** Awareness and casual interaction information is especially time sensitive and must be kept up to date. Users also require knowledge of their current state in terms of past actions and future goals for planning.
- 9 **Provide methods for controlling focus.** As an individual's interest in their locales changes over time, he or she needs to be able to adjust his or her focus onto the locales as well as the people and artefacts within them.
- 10 **Provide methods for controlling nimbus.** Similarly, an individual needs to be able to adjust his or her nimbus (what others are able to see). That is, people need to adjust how they are visible within their interaction context, or to restrict what others can see because of privacy concerns.
- 11 **Represent changes in awareness by varying information content and interaction affordances.** The user's awareness of another person, artefact, or place is represented by the information that is available about it. When the awareness is changed, either through a variation in the focus or the nimbus, the presented information content should change accordingly.

As simplified restatements of a rich previous work, these principles omit many of their subtleties if taken by themselves; thus we recommend that designers who use these should be at least somewhat knowledgeable of the underlying theory. We also believe that the richness of these principles lies in how they can be combined vs. considering each in isolation. Our own design of the Community Bar, described in Section 4, stems from four particular combinations of the principles, outlined below, that shaped our thinking of systems supporting casual interaction.

Multiple social worlds through public locales. Principles 3, 5 and 6 argue that each person inhabits multiple social worlds, which means that any system should somehow let a person inhabit these multiple worlds at the same time. Principles 3, 4, 5 and 8 talk about the richness of the information shared by these groups, which we believe implies that the people, artefacts, and conversations that define a locale should be publicly visible to all who are part of that group, at least by default. The three main categories of tools that support casual communication between distance-separated collaborators – Instant Messengers [12], chat rooms [4], and video-based media spaces [1] – have strengths and weaknesses that can be related to the principles. IM performs well with regard to Principle 2, for it allows lightweight transitions from awareness to interaction. However, IM is most notable for its lack of support for group interactions as specified in Principle 3. That is, it does not readily support small groups of intimate

collaborators. IM communication is oriented to communication between a pair of individuals and, while group conversations are possible, they require explicit setup and so lose many of the benefits of true group-based interaction, e.g., overheard conversation and broadcast queries [6] and easy entry into group conversation [4][8][9][10].

While chat rooms and their variants do support group interaction, they lack the features that make IM so successful, i.e., providing informal awareness (Principle 4), and the means to easily and casually enter into interaction (Principle 2). For example, TeamWave Workplace [8] offered a community with media-rich rooms supporting functionally powerful groupware interaction. However, its implementers noted that its main failing was that people could not easily tell if others were in a room; thus opportunities for real time interaction were lost.

Always-on video media spaces (VMS) are IM-like in that they provide very rich awareness through the video image, and easy means to move into casual interaction over that same channel. However, they vary greatly in how causal interactions are supported i.e., most favour diads (i.e., point to point full video/audio connections), while a few use picture in picture technology to support somewhat larger groups. Very few are like chat rooms, where all conversations are seen; those that do so are mostly variants of meeting room technologies. Notable exceptions include the Notification Collage (NC) [9][15] (a predecessor of our Community Bar) that embeds snapshot video of participants in a publicly viewable space, and that lets people easily enter into chat and other media-rich communications. However, to our knowledge, no VMS provides facilities for multiple locales (Principle 5). They are rarely peripheral (Principle 1), usually consuming a large amount of screen real estate. Most have no means to adjust focus or nimbus (Principles 9-11), except in very rudimentary ways (but see [13] for an exception).

To foreshadow what is to come, the Community Bar extends our earlier work in the Notification Collage project [9][15] by adding facilities for multiple locales, by presenting the information peripherally, by having all information within locales publicly viewable by the social world, and adding focus / nimbus controls.

Ad hoc groups. The research on casual interaction and the Locales Framework (Section 2, and Principles 3, 5-8) tell us that social worlds are plentiful, that their membership (and member involvement) fluctuates, that they may interrelate to one another, and that they have different lifetimes (some are long standing, some form and dissolve rapidly). Yet most systems have very rudimentary notions of ‘groups’, let alone the richer concept of a social world. While buddy lists in Instant Messengers can be created and adjusted by an end user, these are contact lists, not groups. Most media spaces are designed around a single group or community; if ad hoc groups are formed between members, this happens in spite of the technology rather than because of it. Chat rooms are usually centred on special interests and/or topics; while a social world can form around a chat room, it takes a while for its identity to be formed and its culture to develop.

Clearly, a limitation in most these systems is that the formation of ad hoc groups is unsupported or heavyweight. In the Community Bar design, a priority was to support quick, lightweight formation of social worlds by enabling people to rapidly setup and join both short and long-term locales.

Lightweight transitions from peripheral awareness to foreground interaction. Various principles collectively suggest a tension between a person’s desire for a minimal amount of unobtrusive yet dynamic awareness information of their intimate collaborators (Principles 1 and 3), the need to acquire and explore richer forms of that information or to open rich communication channels as desired (Principle 4), and the need to act upon that information and/or engage in communication (Principle 2). As we will see in the Community Bar, we relieve this tension by offering people a progressive view of information. Basic awareness information is placed at the periphery of their screen, and they drill down into that information to gain content and to engage in conversations.

Focus/nimbus control of centre/periphery relationships and mutuality. Finally, most IM, chat room and media space systems give people a binary choice for their involvement. They are either in or out as a member, and they either see or don’t see awareness information (but see [2] for an exception). This kind of design is directly contradicted by our principles. In particular, we believe the centre and periphery relationship – a central tenet of the Locale Framework – can be represented by the focus/nimbus model (Principles 9-11). Mutuality is the most obvious application of focus/nimbus and the context in which it has been described in this paper. However, the model also applies to other aspects: membership of locales, individual views and viewsets, and interaction trajectory.

Again foreshadowing what is to come, people using the Community Bar can express their involvement with a locale by adjusting both their nimbus and focus. That is, they can afford or restrict what others see, and can also adjust their viewable details of the locale and its artefacts. In this way, membership of the locale is not binary. By affording or restricting focus/nimbus, membership becomes a fluid movement from centre to periphery.

In summary, the above discussion illustrates the collective implications of the 11 general design principles. No system uses them all: e.g., while Orbit [11] uses Locales, it focuses on work processes rather than how to transition into them. We used them to motivate our design of the Community Bar (CB), discussed next.

4. COMMUNITY BAR

Community Bar (CB) is a groupware tool that provides the ad hoc and long term groups comprising social worlds with informal awareness information of group members, their conversations and their artefacts, which in turn leads to further casual interactions. Its major design decisions were derived from the above-mentioned principles. ‘Smaller’ interface decisions were heavily influenced by the continuous feedback provided by an evolving community of about 20 HCI researchers in our laboratory who were using CB continuously for over a year, and who had used the Notification Collage [9] – CB’s predecessor – for the preceding four years. As well, the presentation of information within CB is heavily inspired by Microsoft Sideshow’s [3] sidebar; it locates basic awareness information at the periphery while allowing quick drill-down into information when people want to move from awareness into interaction. To this, CB adds the idea of groupware media items found in the Notification Collage [15][9]. These become the building blocks of media-rich locales. In the rest of this section, we describe the basic user

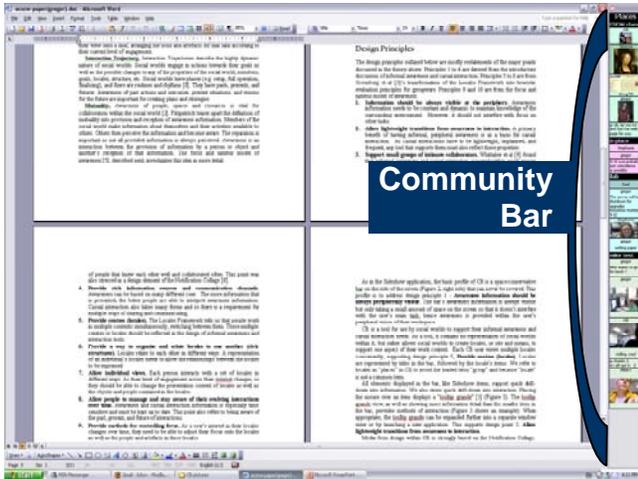


Figure 2: Full screen capture; Community Bar on the side

experience of CB, and then detail and relate our design choices to the principles raised earlier.

An Overview of CB. The basic profile of the CB is a space-conservative bar on the side of the screen (Figure 2, right side; shown in full in Figure 3) that can never be covered by other applications [3]. By a popup menu available through the ‘Places’ tile at the top of the bar, people can create and/or join a ‘Place’ within CB; each place is a locale that defines a site and means for a social world. For example, Figure 3 shows an individual’s viewset of four online places, titled: CSCW class, G-place, ilab, and mike test. Each place contains a number of multimedia items [15][9], representing things like people in the world (as live video, photos or names), public conversations (as chat dialogues or sticky notes), or group artefacts (e.g., web pages of common interest). The place names, the membership of people to that place, the choice of media items within them and the content of these items are completely defined by the group on a moment by moment basis. All members of a place will see something similar to Figure 3, although each member’s viewset of places and how they view particular media items may differ.

When a media item captures the attention of a person, he or she can explore and even interact with that information in greater detail. First, when the person places the mouse over an item in the bar, CB displays a “tooltip grande” [3]. The pop-up in Figure 3 (left side) and in Figure 4a (left) illustrates how this works for the Presence media item. Mousing over the sidebar’s low fidelity and infrequently updated video image of Gregor raises its tooltip grande, which contains a higher fidelity and more frequently updated image as well as various controls. When that person clicks on the title bar of the tooltip grande, a new separate window called the “full view” displays even richer information, and makes available all the functional capabilities of the item. This view may vary depending on who is looking at it. For example, the full view of Gregor’s Presence item, as seen by people other than Gregor, is shown in Figure 4a, right. It contains even higher resolution and higher frame rate video, his picture, and offers its viewer the ability to enter into a vocal conversation through the ‘Push to Talk’ button. Gregor sees this view somewhat differently (Figure 4b, right), where it offers him controls on how to change how others see him, e.g., as a photo or as a video.

- P = Presence Item
- C = Chat Item
- S = Sticky Note Item
- W = Web Item

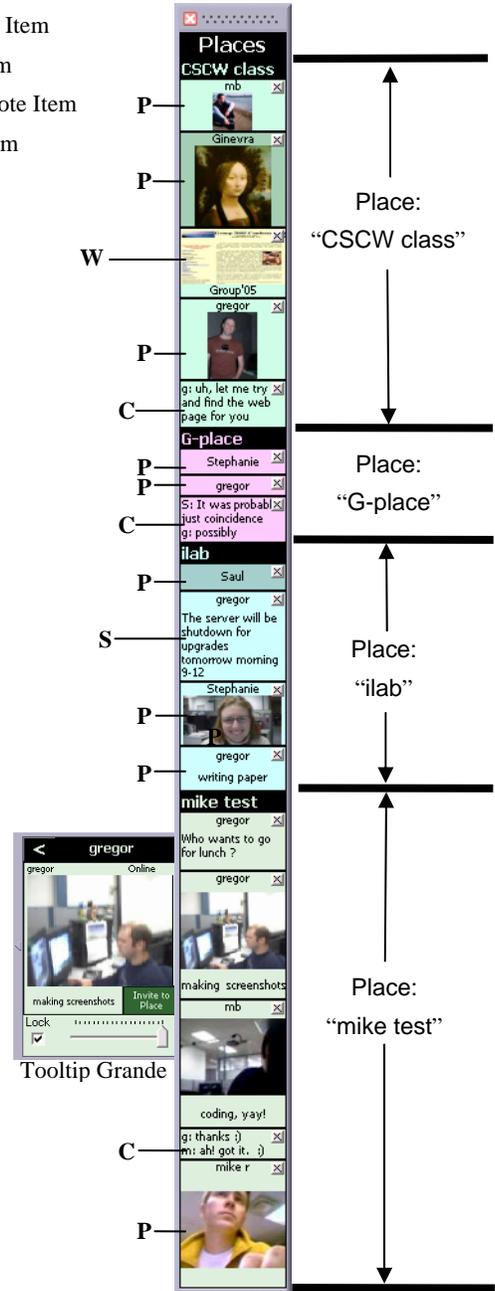


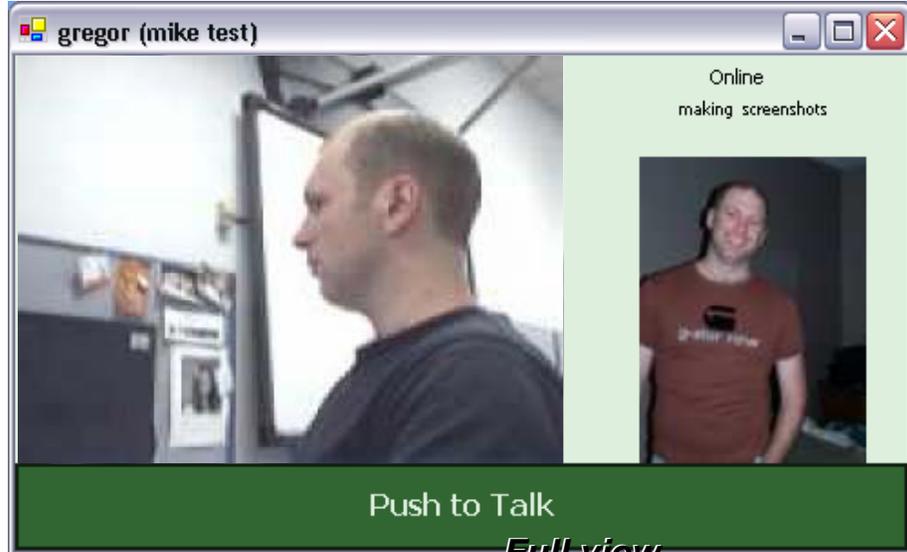
Figure 3: Community Bar. Visible are four labeled places, 4 types of items, and the presence tooltip grande.

Similar capabilities exist for other media items. For example, the tiled, tooltip grande and full view of the chat dialog item are visible in Figure 6. Of special note is the full view of a Place, which fits all the tooltip grande view of a place’s media items into a window as a rectangular grid (not shown). In this manner, the full view of a Place almost completely implements and therefore subsumes all capabilities of the Notification Collage [15][9].

All tooltip grande popups contain a ‘focus’ slider control (e.g., as seen in Gregor’s tooltip grande in Figure 4a) that allows the user to control their personal view of items. Moving the slider from right to left not only shrinks the media item’s size in the bar, but it also semantically changes the information so that it is appropriate



Tooltip Grande

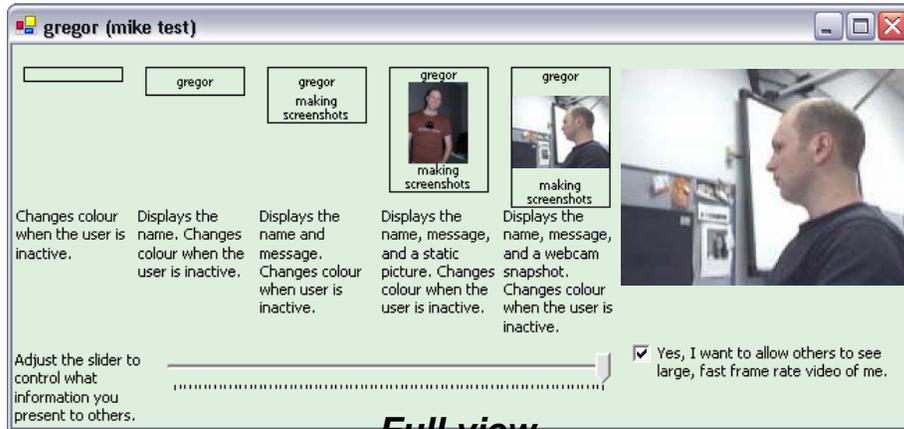


Full view

A. Other participants' view of the Presence media item



Tooltip Grande



Full view

B. Owner's view of the Presence media item



C. Progression of the Presence Tile (in the bar)

Figure 4: Various views of the presence item. The top / middle show the tooltip grande and full views for the participants other than the owner / the owner respectively. The bottom shows different awareness levels of the Tile view.

to its reduced size. Similarly, the 'owner' of the media item (i.e., the person that created it) can adjust a 'nimbus' slider control in the separate window view (Figure 4b, in the full view) to specify a level of detail which others can only see up to but not beyond. Other 'viewers' can personalize this view by using their focus slider to reduce this information even further.

Multiple social worlds through public locales. CB contains no explicit representation of social worlds. Rather, it allows the groups to define and join into these social worlds by creating locales – the *site* and *means* – to support that world's interaction context. These *sites* are the 'Places' within CB, while the *means* are the media items within them and their contents.

CB supports multiple locales (Principle 5), as clearly shown in Figure 3: this particular individual's client displays four Places concurrently. Individuals will have their own particular sets of Places visible on the display; this comprises his or her individual viewset. The inter-related aspects of these worlds (if any) also define their civic structure.

Each Place comprises different sets of media items, and thus offers each social world with its own distinctive means. Through these items, people can present themselves to others, engage in conversation, and interact with group artefacts as desired. Each person can act in distinct ways in each of the Places they inhabit.

Within a Place, all information and interactions are public to every person in the Place, supporting true group mutuality and interaction (Principle 3). In this sense we provide an interaction style similar to that of chat rooms. Locale members are able to share awareness information, to send broadcast queries (e.g. “Is there anyone who knows about X?”), to overhear conversations and join those of interest to them.

Ad hoc groups. CB supports quick, lightweight formation of social worlds by letting people rapidly set up both short and long-term Places, to join and populate them, and to add information and activities to them that become their ‘means’. Unlike most groupware systems, this means that ad hoc means groups are well supported.

New places can be created and named by any person at any time. As shown in Figure 5, this is done easily by form-filling within the ‘Places’ tooltip grande. Once created, anyone can join or leave a Place by selecting or deselecting it from the offerings in the list. Places persist until the last member leaves, although an option exists to keep a Place around even though no one is part of it. Place membership persists across sessions; when a person is not logged on, she is listed to others as an offline participant (e.g., see Figure 6, right side). This accords with membership in a social worlds, for people will belong to that world (albeit at the periphery) even when they are not present.

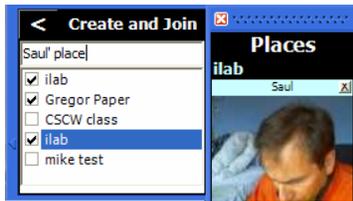


Figure 5: Places tooltip grande for creating and joining Places.

People can invite others to locales, by selecting the ‘Invite to Place’ button in the Presence item’s tooltip grande (see tooltip grande in Figures 3 and 4a). The invitation shows up as an item on the invitee’s bar, and the tooltip grande provides an option to accept, in which case they automatically join the locale.

These three facilities – creating, joining, and inviting – provide support for not only long-standing social worlds but ad hoc groups as well. For example, let us say two people in an existing place start a conversation, where they realize that this should best be continued apart from the current social world. They may quickly create a new Place and join this. One person may then invite a third person in from another Place. Others may populate this place with media items and information relevant to the purpose of this world. Depending on the whims of the group, this Place may exist for only a short time, or could have long standing existence. The effort involved in this scenario is similar to IM systems that allow more people to be invited into a conversation. However, unlike IM the social world that is created has a lifetime that can go well beyond the momentary conversation, and it can be populated with other means relevant to the group (i.e., other media items and whatever information people add to it).

Lightweight transitions from peripheral awareness to foreground casual interactions. We used the Microsoft Sideshow sidebar metaphor as our visual shell because its creators designed it specifically: (a) to show awareness information at the periphery (through the sidebar), and (b) to drill-down and eventually interact with more detailed views of that information

(through the tooltip grande and through what we call the full view) [3].

The items presented in CB display rich yet not overwhelming awareness information, and thus afford mutuality (Principle 4). The views are small and of low fidelity but display enough information for people to maintain a low level of awareness of the presence, activity, and availability of others in the social world. Presence items show low quality video snapshots, static pictures, name, and/or status of other people. Other media items present information about announcements, web pages of interest, events, and conversations within the social world; some of these are visible in Figure 3.

However, the Community Bar’s use of the sidebar differs from Sideshow in significant ways. While Sideshow shows information collected from notification services, the Community Bar’s media items are posted by individuals within a Place; thus CB provides communal awareness of a group’s members, its conversations, and its artefacts. Second, CB not only lets people drill down to view information detail held by that group, but these can ultimately lead into direct public interactions (Principle 2). Third, individual items increase their colour saturation if their contents have changed. Thus a person can tell at a glance if anything in the bar has changed since he or she last looked at it. The item reverts to the normal item colour after its tooltip grande is raised, as this suggests that the person has at least scanned its contents. In summary, unlike Sideshow, all CB media items and their information are fundamentally groupware subsystems that lead naturally from awareness information supply to interaction opportunities within a Place.

For example, consider the Chat media item, and how it compares to those found in IM systems or chat rooms. Any person can create a new chat dialog, by selecting options available in that Place’s tooltip grande (not shown). Figure 6 shows the various views of the chat items. Its sidebar tile view is quite small, and just shows an abbreviated form of the last few conversational statements, prefixed by the initial of the contributor. As discussed in the next section, the sidebar view of this dialog can vary in size, in which case fewer lines will be visible; this is illustrated in Figure 6, left. As people converse, the saturation of the chat item’s background colour increases; this indicates that the conversation is progressing (not illustrated). The person raises the tooltip grande version of the media item (Figure 6, middle) to see somewhat more of the now-scrollable conversation, and the full name of the contributors. Perhaps most importantly, they can now add to the conversation simply by typing lines into it. The full view, seen at the right of Figure 6, shows even more information: who is online (i.e., everyone who can see the conversation) and offline (who may see the conversation at a later time if the item is still around), who is typing the message, more of the conversation (which is now formatted using white space), and a larger area for text entry. Unlike the tooltip grande, this view stays on the display until it is explicitly dismissed by the end user. This final view is usually the only one available in chat rooms or IM systems.

All other media items work in a similar manner. Collectively, this progression of views and how they are situated and drilled down from the sidebar allows the user to quickly stay aware of peripheral information, and to easily move into interaction with

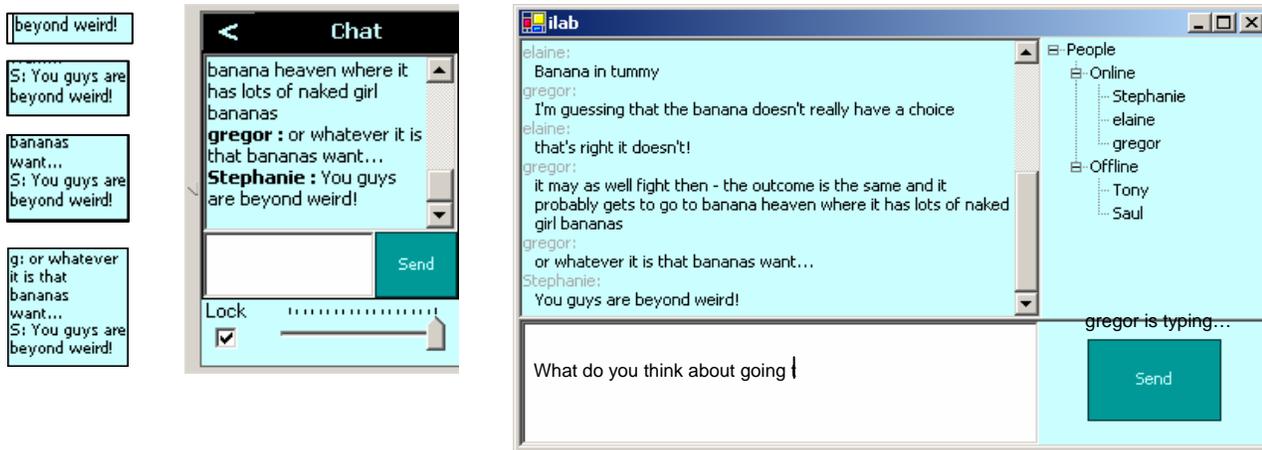


Figure 6: Views of the chat item: Various tile sizes (in the bar), tooltip grande, and full view.

information and people of interest at the time. The selection of view is part of their ability to adjust their individual view.

Focus/nimbus control of centre/periphery relationships and mutuality. If one considers the previous discussion from a focus/nimbus perspective, it should be clear that a person is actually increasing their focus (and their awareness/mutuality) as they drill down through the various media item views, from the sidebar to the tooltip grande to the full view.

However, CB offers other means to give a person explicit control over one's focus on the items and places they are viewing, and on one's nimbus of selected items they have created. By adjusting their focus, they can reduce or expand their awareness of information within a place. By adjusting their nimbus, people can alter how they project awareness information of themselves to others; in turn, this gives the group feedback of how a person situates oneself on the center/periphery spectrum of a particular social world. Focus and nimbus interact to produce the actual view a person sees on the sidebar.

Each tooltip grande contains a slider control (visible in Figures 3, 4, and 5) for adjusting the focus, while the full view may have a similar nimbus control (e.g., Figure 4b). We will use the Presence item to illustrate how this works. As the Presence item is a direct representation of a person, it makes sense for the owner of the item to have control over how they present themselves in the locale. The right image in Figure 4b shows the full view as seen by the owner that includes a nimbus slider at the bottom. The full view also contains a checkbox marked "Yes, I want others to be able to see large, fast frame rate video of me." Nimbus control works to limit what others can see. Unchecking the large video checkbox means that other members of the locale are unable to view the large video normally seen in their full view of that person. Moving the nimbus slider control to the left also limits how others are able to view the owner's information. As illustrated in Figure 4c, it reduces it to slower and smaller video, then to a static picture, and then to a name and text message, and then only to the name of the person. The top part of this form in Figure 4b is live feedback on how the view has been restricted – the possible views are blacked out as they are removed.

Viewers of this item can see up to – but no more than – the amount of information allowed by this nimbus projection. However, a viewer can further reduce the fidelity of this nimbus

information as well as the space taken up by it by moving the focus slider found in the tooltip grande, visible in Figure 4a. Internally, CB combines the focus and nimbus values for a particular item to calculate an 'awareness value', which controls the size that the item has available to display itself in the bar. The item then displays as much information as it can in the space available. As seen in the various bar views of the Presence item in Figure 4c, when there is very little space the item is blank and only shows online/offline status by changing its colour saturation level. As it is given more space it is able to show a name, then a text message, then a static picture, and finally a small video (assuming the nimbus value allows it). A similar effect is shown with the chat item at the left of Figure 6.

CB also automatically reduces the focus on items when there are too many items to fit on the side of the screen. Its compromise is to calculate awareness values where all items just fit in the display. However, a person can adjust this further through the focus slider provided on each Place's tooltip grande. For example, if they are more interested (i.e., more in the center) of one place than another, they can increase their focus on it; the focus on all the items within that Place are uniformly increased. This automatically reduces a person's focus on other places when screen space is at a premium, i.e., it moves the person more to the periphery of those places until all elements fit on the display.

As mentioned, the full view of a Place contains the tooltip grande view of all media items, on permanent display. This is another means for people to increase their focus onto one or more places.

Providing control over focus and nimbus in these simple, lightweight mechanisms – which are actually much easier to use than this wordy description would suggest – allows each person to manage their mutuality with other people, to fluidly control their centre/periphery membership in each locale, to change their engagement with artefacts, and to adjust their individual views. They do this dynamically as their needs and interactions evolve.

5. DISCUSSION

The strength of Community Bar lies in two aspects: firstly, as we have been discussing to this point, it is firmly based on a unification of theoretical principles from a number of different and well-established research sources; and secondly, design details are based on a year of participatory design and use of

Community Bar (CB) from a community of users as well as a further four years of participatory design with its predecessor, the Notification Collage (NC). In this section we discuss the historical context and evolution of the NC (given in detail in [15][9]) and the CB in the context of how participatory design has shaped and continues to shape the design and research.

The first incarnation of the NC was a public electronic corkboard displayed on a large touch-sensitive and wall-mounted Smart Board, designed for use primarily by a collocated group. People were able to post information to the board from their desktop computers, such as text announcements and video, for other members of the group to see. Because viewers and posters were collocated, it augmented rather than replaced normal face to face communication. NC use increased significantly when it was extended to be available over the Internet; distant people were able to post to it, and could view an instance of the display on their own desktops. Telecommuters (mostly those working from home) and travelling members of the laboratory (e.g., those on internships or extended research visits) used it not only to stay aware and interact with the rest of the group, but to feel socially involved with them during periods of absence. People modified their use of media items to reflect this; e.g., the NC sticky note item was intended to post announcements, yet people co-opted it to afford live public group discussion. The NC investigators responded to these new uses, as well as to suggestions from the NC community, by altering the system to reflect group desires. Consequently, newer generations of the NC had integrated client / viewers running on desktop machines, and included a rich (and extensible) set of different media items supporting awareness and interaction. Due to this interactive and participatory design with the user community, the original vision of an electronic whiteboard for collocated groups evolved into a virtual public room useful to *both* collocated and distributed groups.

As the NC user community grew, an increasing body of user comments and observations made clear some limitations of the fundamental design. Most notable amongst these were: (1) it consumed too much screen real estate; thus the people who used it were mostly limited to those with two monitors; (2) subgroups were forced to work in the single public space of the NC; they sometimes chose to switch to other communication means when this was inappropriate; and (3) presence was oriented around video snapshots, alienating the camera-shy and those without cameras. These problems were serious enough to discourage a subset of community members from adopting NC.

The first version of Community Bar was designed to address these issues. Screen real estate problems suggested use of the SideShow [3] sidebar containing resizable tiles. Sub-groups were supported via multiple Places. Alternatives to video-based presence items were created. In later versions, we worked on the derivation and unification of the theoretical principles discussed earlier to provide a design structure for the radical changes that we were making. The Locales framework enforced the need for multiple ad hoc groups, where we crafted the ‘site and means’ that these groups should have available to them. The Focus/Nimbus model suggested that people could adjust what they wished to reveal to others in a single Presence media item by adjusting their nimbus.

Currently, Community Bar is in everyday use throughout our research community. It is regularly used by collocated people, by telecommuters, and by members currently residing in other cities

(e.g., Vancouver, Boston, Redmond). Some evidence for the success of our changes is that individuals that didn’t use the Notification Collage are now using CB. However, Community Bar is a living research project and is still evolving and changing. We are continuing our pattern of participatory design, responding both to feedback and observation of use. We are also planning more formal evaluation with a larger, more disparate community.

Multiple social worlds through public locales. The current user community uses CB in an IM-like fashion to maintain awareness and to easily move into interaction with both collocated and distributed group members. Yet they also use it in a Chat room fashion as conversations are overheard, broadcast queries are issued, and artefacts such as web pages are shared with the group. As hoped, CB does support a social world by offering it a public locale (the CB place that is the site), and the means (media items).

However, the multiple Place functionality is not yet used heavily. This likely arises because the group currently using the CB is fairly cohesive, and enjoys working within one large Place. This group does not see a need to splinter themselves into long-term sub-groups. When people who did not know the main user community joined CB, they created and joined Places other than the main Place, suggesting that the large single Place tends to come from this particular group and their desire to maintain a single social world. This is an acknowledged bias in our user community that we are seeking to address now that CB is ready for wider distribution. As the community increases and becomes more heterogeneous (i.e., as multiple social worlds are present), we expect that we will see them create both short and long term multiple locales, allowing us better insight into how to best support multiple social worlds and the civic structures (principle 6) that arise from them.

Ad hoc groups. Even within this single social world, we see that people create temporary Places for ad hoc sub-groups working on intense, well-defined projects or activities. These groups typically move their specialized discussions and activities (which may not be of interest to the larger group) to another place, because it provides the sub-group with a focus, and because it avoids distracting the larger group. One example is paper writing, where co-authors create a Place that mediates their conversations about the paper. Another example is training. When a new person is invited into an existing place, an online group member will typically create a new place to ‘train’ that other person.

The lifetimes of these ad hoc group Places have generally been on the order of a few hours to a few days. In contrast, the main Place has a very long lifetime, spanning almost a year.

Lightweight transitions from peripheral awareness to foreground interaction. Current feedback and observation indicates that Place members easily move from awareness to interaction, either by conversing with present group members or by exploring information artefacts left by others. That is, the CB (and the sidebar metaphor [3]) fits well with people’s desired work practices. People will glance at the bar to maintain awareness, use the tooltip grande to add quick comments or to see information details, and open the full view for extended interactions. The fundamental design operates well; CB design iterations have concentrated on small interface refinements desired by the group.

Focus/nimbus control of centre/periphery relationships and mutuality. The focus and nimbus features are also well liked and well used. People commonly adjust their nimbus to fit different work settings. For example, some telecommuters reduce nimbus to substitute a photo for video, because they may not feel video is appropriate to their home setting. People also adjust their focus, mostly to reduce the view of items that have already been viewed and that are not changing. Yet we recognize that this focus/nimbus control can be interface overhead; we are now investigating automatic adjustment of focus and nimbus. For example, should CB automatically reduce a person's nimbus when they have been idle for a while? Should CB reduce a person's focus on information if that information has not changed? Early experimentation suggests that such mechanisms, while having potential to be useful, have to be very carefully designed to avoid distraction.

6. CONCLUSION

Existing informal awareness and casual interaction tools, while demonstrating the benefits of some design points, are impoverished when considered in the context of a rich sociological concept of community. To address this, we designed the Community Bar using both a comprehensive sociological theory and the input of five years of participatory design. Our reflection on the theory behind CB, its design, and its use has led to two primary contributions.

Our first contribution is to derive a series of design principles from a comprehensive sociologically based framework, supported by a generalised model of awareness and previous research in informal awareness and casual interaction specifically. While we do not claim that individual principles are new, their combined effect gives new insights into the design of systems, as discussed in Section 3. These combined principles indicate support for multiple public locales, ad hoc groups, lightweight transitions from awareness to interaction, and focus/nimbus controls for membership and awareness.

However, we recognize that these principles are high-level and very general, and do not immediately translate into a particular design. Consequently, our second contribution is to demonstrate how these principles can be applied to generate a novel design of groupware supporting informal awareness and casual interaction: the Community Bar. Of course, even good principles do not necessarily guarantee good design; this is why CB design details are augmented by five years of observation and feedback: four years from its Notification Collage predecessor, and one year of CB use (this includes a pilot study that tracked details of user interactions within CB).

Future directions for Community Bar involve expanding its offerings of media items. More importantly, we are just about to deploy the system to different communities; evaluation of its adoption and use by these communities will add to our knowledge of how social worlds should be supported.

7. ACKNOWLEDGMENTS

Research funded by NSERC NECTAR Research Network. Special thanks to the Interactions Laboratory at the University of Calgary who provided feedback on CB use. Michael Boyle,

Michael Rounding, and Carman Neustaedter helped on earlier drafts of this paper.

8. REFERENCES

- [1] Bly, S.A., Harrison, S.R., and Irwin S. Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment, in *Comm. ACM*, 3, 1, (1993), 28-47.
- [2] Bradner, E., Kellogg, W., & Erickson, T. The Adoption and Use of Babble: A Field Study of Chat in the Workplace. *Proc ECSCW*, (1999), 139-158.
- [3] Cadiz, JJ, Venolia, G.D., Jancke, G., and Gupta, A. Designing and deploying an information awareness interface. *Proc ACM CSCW* (2002), 314-323.
- [4] Curtis, P., Nichols, D. A. MUDs Grow Up: Social Virtual Reality in the Real World. *Proc 39th IEEE COMPCON* (1994), 193-200.
- [5] Fitzpatrick, G. *The Locales Framework: Understanding and Designing for Wicked Problems*. Kluwer Academic Publishers, (2003).
- [6] Fitzpatrick, G., Parsowith, S., Segall, B., Kaplan, S. Tickertape: awareness in a single line. *Proc ACM CHI* (1998) 281-282.
- [7] Greenberg, S., Fitzpatrick, G., Gutwin, C., and Kaplan, S. Adapting the Locales Framework for Heuristic Evaluation of Groupware. *Australian Journal of Information Systems*, 7, 2 (May 2000), 102-108.
- [8] Greenberg S. and Roseman, M. Using a Room Metaphor to Ease Transitions in Groupware. In M. Ackerman, V. Pipek, V. Wulf (Eds) *Sharing Expertise: Beyond Knowledge Management*, 203-256, Cambridge, MA, MIT Press. (2003)
- [9] Greenberg, S. and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. *Proc ACM CHI*, (2001), 515-521.
- [10] Kraut, R., Egidio, C., Galegher, J. Patterns of Contact and Communication in Scientific Research Collaboration. In *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Lawrence Erlbaum Associates Publishers, 1990, 149-181.
- [11] Mansfield, T., Kaplan, S., Fitzpatrick, G., Phelps, T. Fitzpatrick, M. Taylor, R. Evolving Orbit: a process report on building locales. *Proc ACM Group*, (1997) 241-250.
- [12] Nardi, B.A., Whittaker, S., and Bradner, E. Interaction and Outraction: Instant Messaging in Action, *Proc ACM CSCW* (2000), 79-89.
- [13] Reynard, G., Benford, S., Greenhalgh, C. and Heath, C. Awareness Driven Video Quality of Service in Collaborative Virtual Environments, *Proc. ACM CHI*, 1998.
- [14] Rodden, T. Populating the Application: A Model of Awareness for Cooperative Applications. *Proc. ACM CHI*, 1996, 88-96.
- [15] Rounding, M. *Informal Awareness and Casual Interaction with the Notification Collage*. M.Sc. Thesis, University of Calgary, Alberta, CANADA, 2004.
- [16] Whittaker, S., Frolich, D., and Daly-Jones, O. Informal workplace communication: What is it like and how might we support it? *Proc ACM CSCW*, (1994).131-138