A Field Study of Community Bar: (Mis)-matches between Theory and Practice

Natalia Romero
Eindhoven University of Technology
Den Dolech 2, 5600MB
Eindhoven, The Netherlands
+31-40-247-8398
n.a.romero@tue.nl

Gregor McEwan
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2N 1N4
+1-403-210-9499
mcewan@cpsc.ucalgary.ca

Saul Greenberg
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2N 1N4
+1-403-220-6087
saul@cpsc.ucalgary.ca

ABSTRACT
Community Bar (CB) is groupware supporting informal awareness and casual interaction. CB's design was derived from three sources: prior empirical research findings concerning informal awareness and casual interaction, a comprehensive sociological theory called the Locales Framework, and the Focus/Nimbus model of awareness. We conducted an in-depth field study of a group's on-going use of Community Bar over several weeks. We use results obtained from this study to reflect upon the matches and mis-matches that occurred between the theoretical usage behaviour predicted by our theoretical design principles versus the actual usage behaviours observed in the deployed implementation. As a critique, this reflection is an important iterative step in considering how CB should be redesigned, and serves as a cautionary tale of the difficulty of translating theoretical nuances into practice.

Categories and Subject Descriptors
K.4.3 [**Computers and Society**]: Organisational Impacts – *Computer-supported collaborative work*.

General Terms
Design, Experimentation, Human Factors, Theory.

Keywords
Locales, casual interaction, distributed 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 [11][18]. 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 [11]. In contrast to formal meetings, casual interactions are unplanned, brief, frequent, and usually engage small groups of people familiar with one another [18]. 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 [11][18]. 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 [11][18]. 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) [14], chat rooms / MUDS [5], 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 a member 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 long-term 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 both empirical studies of casual interaction behaviours [11][18] and social science theory [6] [16]. In particular, we are motivated by the Locales Framework [6], one of the few comprehensive theoretical group interaction frameworks in the computer science field, as well as the Focus and Nimbus model of awareness [16]. We have previously derived and combined tenets from these studies and theories into a set of design principles [12]; these are

Cite as:
Romero, N., McEwan, G. and Greenberg, S. (2006) A Field Study of Community Bar: (Mis)-matches between Theory and Practice. Report 2006-826-19, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, T2N 1N4. March 17th.

Table 1. Principles and descriptions

1. Awareness information should always be visible at the periphery	Informal awareness information is constant and dynamic yet in the background. It should always be displayed but not distract from foreground tasks.
2. Allow lightweight transitions from awareness to interaction	When awareness information change captures the user's interest, they must be able to easily investigate and transition to more detail and interaction opportunities.
3. Support small groups of intimate collaborators	Casual interaction occurs between intimate collaborators, who know each other and work and interact together often.
4. Provide rich information sources and communication channels	People gain awareness and interact through incredibly rich channels in the physical world and these abilities can be better leveraged by rich virtual channels.
5. Provide centres (locales)	Groups make use of locations, tools and artifacts to during their interactions and so a virtual tool must also provide these things to support awareness and interaction.
6. Provide a way to organize and relate locales to one another	Groups have relationships between them that affect their interactions and so the system must represent the important relationships.
7. Allow individual views	Individuals are members of different groups and have unique views within each group. The system must allow individual views within the group and across groups.
8. Allow people to stay aware of evolving interactions over time	Historical information and future plans are vital for people to form goals and manage their interaction.
9. Provide methods for controlling focus	Each person has unique requirements for their view onto other people, locations, tools and artifacts. The system needs to give control over these views.
10 Provide methods for controlling nimbus	People need to be able to manage their appearance and identity as presented to others.

summarized in Table 1 but not described further due to lack of space. We then used these principles to design the Community Bar (CB), a groupware tool that supplies ad hoc groups with rich awareness information leading to casual interaction [12].

In this paper, we present a field study of Community Bar's use by a small group. We use results obtained from this study to reflect upon the matches and mis-matches that occurred between the theoretical usage behaviour predicted by our theoretical design principles versus the actual usage behaviours observed in the deployed implementation. As a critique, this reflection serves not only as an important iterative step in considering how CB should be redesigned, but also as a cautionary tale of the difficulty of translating theoretical nuances into practice.

2. COMMUNITY BAR

Community Bar is fully described in [12]. This is just a summary of its key interface components.

Figure 1 illustrates a screen snapshot of CB in use. CB presents itself as a sidebar peripheral display [4], divided into *Places*. Each place represents a sub-group, their communication, their tools, and public information posted to it. Using a menu, people can easily create or join one or more Places. The intention is that each Place serves as a local [6] offering its inhabitants the site and means for group awareness and communication. For example, Figure 1 shows an individual's view of four Places: CSCW class, G-place, ilab, and mike test.

Each place contains a number of multimedia items [17], representing things like people (as live video, photos or names), public conversations (as chat dialogues or sticky notes), or group artefacts (e.g., photos and 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. Multimedia items have three different forms: the *tile view* that shows awareness information, the *tooltip grande* showing more detailed information and controls for minimal interaction, and the *full view* showing all information and interaction possibilities.

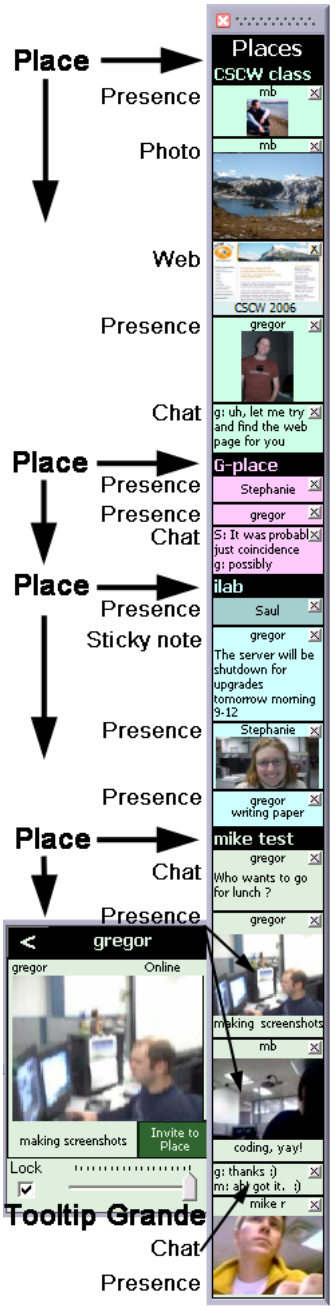

Figure 1: Community Bar

Tiles, meant for peripheral awareness, are always visible in the sidebar. Thus all members within a place will see at least those tiles. For example, all people currently see Gregor's Presence tile, which at this moment is displaying a low fidelity and infrequently updated video of him and text describing his activity. Figure 1 also shows 5 tile types representing people, conversations, or shared information.

When the contents of a tile captures a person's attention, he or she can explore and even interact with that information in greater detail. First, when the person mouses over an item in the bar, CB displays a transient *tooltip grande* [4]. For example, Figure 1 illustrates the tooltip grande for Gregor's Presence media item, which contains a higher fidelity and more frequently updated video image as well as various controls. Second, 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 (not shown). 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, contains even higher resolution and higher frame rate video, his picture, and offers its viewer the ability to enter into a vocal conversation through a 'Push to Talk' button. Gregor sees this view somewhat differently, where it offers him controls on how to change how others see him, e.g., as a photo or as a video.

Similar capabilities exist for other media items. For example, in the Chat Item, the tile view shows the last message or two. The tooltip grande view shows the last 10 messages and allows sending messages. The full view adds to this by showing all messages, the place members, and who is currently typing. Of special note is the full view of a Place, which fits all the tooltip grande views 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 (NC) [17][10].

All tooltip grandes contain a 'focus' slider control (e.g., as seen in Gregor's tooltip grande in Figure 1) that allows the user to control their personal view of items [16]. 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 to its reduced size. Similarly, the 'owner' of a presence media item (i.e., the person that created it) can adjust a 'nimbus' slider control in the separate window 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.

3. FIELD STUDY OF CB IN PRACTICE

Unlike task-oriented productivity tools, Community Bar is intended to support ongoing collaborative social practices as they occur in the everyday world. Consequently, we felt it appropriate to evaluate CB's efficacy through a field study investigating how people used CB while continuing with their normal practices. Yet we recognized that CB's use by a group would evolve over time, where it would be adopted into the group's everyday social practices and their cultural norms. We were more interested in examining how the group used CB after this period of adoption, so that we could see how their social practices had stabilized. This suggests that a longitudinal field study was needed. The catch is that the logistics of seeding a new group with CB and monitoring them for (say) many months until CB was adopted was onerous. Instead, we decided to study the creators of CB and their colleagues - a group that already had being using CB and its predecessor (NC) for a long period of time.

While examining this group introduces some biases (this group is likely more favourable to CB), we stress that the group is still worthy of study for several reasons:

- Members had used CB and its predecessor for several years.
- Current members included people working on quite different projects and who were uninvolved in the actual CB research.
- Group membership has changed over years as people came and went, and thus went beyond the original core group that had vested interests in it.
- All were experienced with casual interaction theory and systems prior to the study, and thus better able to reflect on their practices.
- The principle investigator of this field study was *not* part of this group: her involvement with them was for the express purposes of setting up and conducting the field study.

3.1 Participants

The group we observed consisted of fifteen study participants. All had real world work and social relationships with each other. Eleven were current members a research laboratory at a university (one Professor and ten graduate students), all who saw each other face to face over the course of the work week. Five of the ten graduate students were supervised by the Professor. Three others were former graduate students of this Professor: two had left within the last year and now worked at private industry firms, while the other was now a student at a different university. The final participant was a researcher at another university; this person had a weak academic relationship with this group. Thus fourteen people knew each other very well, while the fifteenth had only met part of the group a few times. Finally, before the start of the study, ten were using CB as a group on a regular basis, four had used it a few times, while one was a new user.

3.2 Method

We collected and analysed activity logs of CB usage, people's subjective diary entries, and performed a series of interviews.

Duration. We monitored all CB interactions between our fifteen study participants for a period of three weeks.

Logging. For the duration of the study CB was instrumented to log all actions both on shared data, such as chat messages and web pages posted, and on personal clients, such as raising full views. This data was then processed and interpreted to obtain information on how CB was actually used.

Diary. A special diary entry media item was created. At any time, CB users could enter stories and reflections about their on-going experiences, which were logged by the system. The diary also prompted people when they did certain CB actions, and when they were affected by other people's actions. For example, if a person adjusted the focus slider, the diary item would ask why.

Interviews. After the study period, eleven participants were interviewed at length about their impressions and experiences with CB. Interview responses were then matched with the logging data and diary item entries to give insight into intentions and experiences as well as the direct actions.

4. Results

We continuously collected 21 days of activity logs. For analysis, we divided each day into four equal parts: morning, afternoon, evening and night. Each of these was then divided into two three hour blocks: morning is 6am-9am, afternoon is 3pm-6pm and so on. Data for three time periods (shown as the gaps Figure 2) was lost due to server problems. Thus a total of 145 time blocks were recorded, representing a total of 435 hours of CB activity.

4.1 Login Activity

Figure 2 graphs the number of people logged onto CB during each block over the whole study period. Of the 145 blocks, only ten were 'empty', i.e., no one was logged onto the system. Not

surprisingly, most of these were during the night (midnight to 9am). For the remaining times, login activity varied in somewhat predictable ways, as revealed by the patterns in the figure. The busiest times were on weekdays, between 9am--6pm. On average, three people were logged in at any one time but during the busy blocks the average rose to six. Peak attendance, up to 10 people, tended to coalesce around the middle of the day.

The Figure also reveals that login activity was not restricted to work hours, as there are still concentrations of people in the evening blocks (6pm-9pm and 9pm-12am). Interestingly, we found that the membership of the daytime group is different from the evening group. To illustrate this, Figure 3 shows one example day where there is a clear change of members (labelled A – J) logged onto the system between afternoon and evening. In this particular case, there is a transitional time overlap, and one member remained throughout the day. This general pattern repeated itself most days.

Further analysis comparing people's actual CB activities in the daytime *vs.* evening 'groups' revealed differences in the content of their interactions. Data collected from chat item content and diary item responses indicates that the daytime group interacted less and that activity tended to be work focussed. In contrast, the evening group typically had more direct interactions that were more playful and socially focused.

4.2 Analysis: Theory *vs.* Practice

The primary purpose of our evaluation was to investigate how the theoretical principles used in the design of Community Bar (reported in [12] and summarized in Table 1) played out in practice. Our analysis below is structured around each principle.

Each of the principles restates a fundamental assumption or property derived from a pre-existing theory. Principles 4.2.1 through 4.2.4 are from informal awareness and casual interaction research, e.g. [11][18]. Principles 4.2.5 through 4.2.8 are from Greenberg et al's restatement of the Locale Framework's principles [6] as groupware heuristics [8], and the remainder are from Rodden's focus/nimbus model of awareness [16]. Each of the subsections briefly summarizes the principle followed by the related experimental results.

4.2.1 Awareness information should be always visible at the periphery.

Informal awareness is, by definition, a background, peripheral process. The information being monitored is constant and dynamic but only gains the attention when the information changes suddenly and sharply. For most of the time, it should not interfere with focus on other tasks. This means that an informal awareness and casual interaction system should constantly display awareness information, but display it in such a way that it is not distracting. CB's sidebar design is heavily based on this premise.

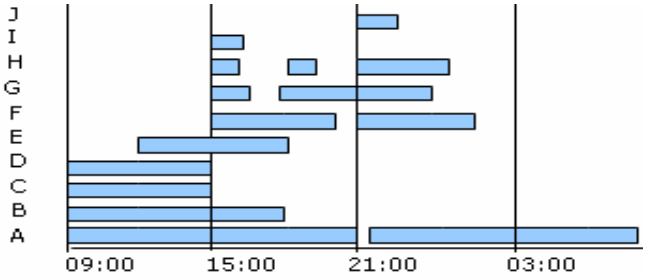

Figure 3: Membership change (note 15:00-18:00 period)

Our study data indicates that CB's design largely matched this principle, where people's primary use of CB was for awareness.

From an interface perspective, participants indicated that the benefit of the awareness information was worth the screen space that they had to sacrifice. A typical response:

"I think in general, the amount of space that it uses for the information it gives off is quite balanced … it's definitely useful for the size that it is."

CB's design tries to trade off peripheral vs foreground awareness through visibility and screen real estate. One person commented on this when comparing CB to its Notification Collage (NC) predecessor. She said that NC, which displayed larger notification on a full-screen, led to more interactions [10][17]:

"Even though I like that CB takes up less space, I probably interacted with [NC] more and used it more when it … took up a whole monitor … I would post more things and I would use more things other than the videos."

People also adjusted the focus (the size) of each tile to reflect their ongoing interest in its information as awareness:

"[I make the tile bigger] so I can see new messages when they arrive"

"[I make the tile smaller] because there is nothing in it now – I will open it up big again if it turns bright blue [which indicates new information]"

Distraction can occur either because the interface itself is problematic, or because information content changes so quickly that it demands their attention. To check this, the diary routinely asked "Are you too distracted by CB at the moment?" In all cases respondents said that (excepting the diary queries) they were not distracted, largely because group activity only placed modest demands for their attention:

"No, as … there is little direct activity other than the video i.e., I am 'up to date'."

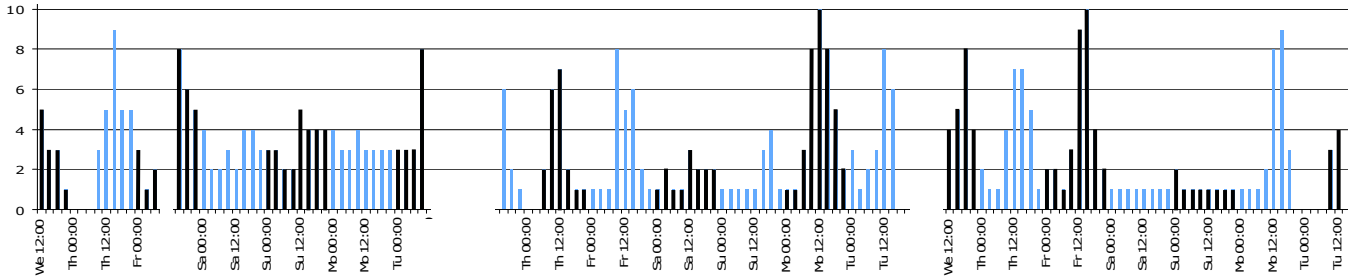

Figure 2: Number of people logged on over the study period. Days in alternating colors, blanks on the X-axis indicate missing data

"I only check every once in a while if something new is going on and it rarely is. It would be different if people were more active on the CB."

An item's salience may provide useful awareness information to one person, and distraction to another. For example, when a new message is posted to a Chat item, the tile subtly notifies people of this change by increasing its colour saturation. Interview responses to the utility of this fell into three categories:

Distraction: "*Sometimes it's too distracting, especially if there's a conversation going on that I'm not interested in."*

Ignoring: "*I would say it wasn't distracting enough if anything because it was really hard for me to tell when people were talking to me."*

Involved: "*I will notice when the text items change and I'll go and look and see if it's a conversation that, you know, I want to participate in and if it's not I'll just go back to what I'm doing."*

The *distraction* response was given by people who felt compelled to look at the changed chat item but who that were uninvolved in its conversation. The *ignoring* response was by people who learned how to ignore this notification, but then felt left out if the conversation turned to something that they should have seen. The final *involved* response was for people who felt directly involved in the conversation. These results suggest a peripheral awareness display cannot satisfy all people at all times, for the balance between useful awareness and distraction is heavily subjective and determined by context.

4.2.2 Allow lightweight transitions from awareness to interaction.

In the real world, awareness creates opportunities for brief but rich interaction [11][18]. CB's drill-down design, adapted from [4], leverages this idea. When awareness information in a tile captures the user's attention with something of interest, the CB participant has the option of using the tooltip grande and/or the full view to drill into more detail and to interact with that information.

Participants responded extremely favourably to this aspect of CB's design, as typified in the following quotes about the chat item:

"I do use the [tooltip] a lot to chat because it's convenient. Mostly for short conversations though, if it looks to be a long conversation then I'll open the separate."

"[I use the chat full view] as an easy way to type as well as view conversation. Also, to see if the other person is still typing."

All of the participants said that the chat item views reflected their desired interaction extremely well. People used the tooltip to see a bit more of the conversation while giving them the ability to send quick messages. People used the separate view for extended interaction and to view the entire conversation.

However, while the overall design principle was praised, people said that some media items did not implement the design principle well. In particular, if a tooltip grande offered only slightly more information and controls than the tile, people did not bother with it. An example is the Presence item: while the tooltip grande video is larger and updated slightly more frequently, it really shows little more than its tile counterpart (Figure 1). A quote concerning the Presence item typifies what most people said:

"I don't think I actually use the [Presence] tooltip … I think all I ever do is expand the tooltip to get at the arrow to open up into the bigger [Full] view where I can then adjust to turn on my video."

In spite of this design problem with some of the media items, we were told that bypassing the tooltip was not a big deal as it was very quick and easy to open the separate view.

In summary, compelling media items are those that make full use of their three views, each showing personally significant new information detail leading to progressively richer interactions.

4.2.3 Support small groups of intimate collaborators

Research into informal awareness and casual interaction suggests that the people involved are most frequently already known to each other, work together and interact with each other often [11][18]. We refer to such groups as *intimate collaborators* [16], and these are CB's target users rather than (say) communities of strangers.

From comments received, our study participants were easily divided into two groups: the Professor and the students he was currently supervising; and everyone else. The Professor and his students often worked together closely and interacted with each other often. These people used CB more often for communication. While everyone else also interacted frequently (outside of CB), they had less work ties to the first group. These different relationships came out in the interview comments: the first group was the "core group" and the others who were "peripheral members". As we will see, this led to a divide in how CB was considered.

Core group members consistently talked about the sense of belonging to the community that CB gave them. One participant talked about times when he was unable to use CB from home due to network problems:

"I really lose out, mostly on this feeling of being connected, that I am still part of the group, especially if I'm working at home because there's a problem… there's no-one else around and it's very isolating."

In contrast, peripheral members often reported that they felt like outsiders, and that most of the explicit communication on CB did not involve them. As one member said:

"As it happens right now I'm not working that closely with anybody in [the main CB location] … so if it were a time when I was working more closely with people I could see where it would have been more useful."

Another, who was collocated with the core group but not part of it, summarized why she felt left out:

"I think most of the conversations are just [the Professor participant] wanting something from his students and I don't really care."

However, all people, whether core or peripheral, expressed sentiments on how useful CB was for maintaining an idea of what was going on with the rest of the group.

We have little insight into those who did not use CB at all but one of the study participants who worked with other groups commented on how useful CB would be in those other groups:

"Do I think that it could be useful? Yes. Do I think that [my own groups] would use it? No. I'm sort of split between two groups

which are theoretically doing the same stuff but are not very cohesive … I think that there needs to be some social cohesion and I think that there needs to be some work cohesion"

In summary, CB works best for small coherent groups of intimate collaborators. It is works less well for participants who are peripheral members. If people are not part of a cohesive group, they do not see CB as a panacea for bringing it together.

4.2.4 Provide rich information sources and communication channels.

In a collocated, physical environment, people use a wide array of incredibly rich awareness and interaction channels; they can see, hear, smell and touch and share artefacts in many subtle and varied ways. Though a technological system like CB cannot hope to mirror the richness of the everyday world, its design favours information-rich over information-poor channels. For example, a picture is richer than a text name, and video (even low frame-rate) is even richer and communicates more about the person. As well, CB is designed as an open-ended system, where new media items can be created and included into it [13]. However, only the media items shown in Figure 1 were available to study participants.

Every participant reported that the rich awareness provided by the video snapshots was their primary motivation for using CB. Typical comments included:

"For me being able to communicate with my colleagues using a tool that is so rich as the CB is very valuable."

"The thing that I like the best is just being able to see video of what everybody's up to … because it's just useful to have a sense of who's in their office and who's on the phone … it just gives you a better sense of what's going on."

Yet the richness provided by the video turned out to be a mixed blessing. Video awareness overshadowed other types of awareness information, e.g., static pictures were considered much less useful. Participants reported that they often neglected other these lesser forms of presence information. One participant, who used video throughout the study, said:

"So the awareness information it gives me [when people are using video] is fairly reliable because I can see when they're around. Whereas if they just have a standard static image it provides me much less awareness. … I don't really notice the away bar … so really I rely heavily on the video. So I really like the fact that lots of people do use the video."

Participants who did not portray themselves through video said they felt somewhat left out. One said that people without cameras were like *"second class citizens."* The relating between video and this feeling of inclusion is emphasised by one person who only started using video in the middle of the study period after he bought a camera:

"There were benefits, in that, you know, people talked to me more … because they knew I was there."

Video presence awareness also caused some frustration because it did not quite give people the capabilities of a true collocated situation. Sometimes people would try to contact someone through CB because they could see that person on the video, but they could not attract their attention. The following anecdote captures this frustration:

"I almost get this impression that [Participant] will only look at the CB every so often, like maybe every 15 minutes or something. I have this impression that he [doesn't notice] when things change on it. It's more of he looks at it every now and then to see if something's happened. …Often there's a long delay when I post a question to him and when he actually finally responds … it bugs me sometimes and I guess I'd like to know why … and I don't get that information."

In summary, CB's richer information and communication channels proved useful for awareness and interaction, but users were still very aware of the difference from real world interaction. This is good motivation to make CB's media items even richer.

4.2.5 Provide centres (locales).

Collaboration involves groups of people working together for a common purpose. The Locales Framework [6] calls these groups *social worlds*. Social worlds make use of locations, called *sites*, and multiple tools, or *means*, to work towards their purpose. The combination of social world with site and means is called a *locale*. As there are many different social worlds needing many different sites and means, technology must provide and manage many locales. Hence CB's Places was designed to emulate multiple locales, including sites and means, for its social worlds.

During the study period, all participants primarily used a single place. Three other places were created, but they were used infrequently and by few people, and were not long-lived. As investigators, we were somewhat surprised by this as we observed many instances where another place would have been warranted, e.g., when a conversation topic diverged or when a subset of the group was working closely together for an extended period of time. However, none of these events resulted in new places being created.

When asked why they did not create new places, participants responded in very similar ways, saying that they were not needed in the existing community social structure:

"In CB you can make multiple places but I've not yet really come across a situation where I need to. Mostly I think because the people who are using CB that I know are all from the same kind of culture, they're all from the lab here … If I'm on the Community Bar, basically what I'm saying I take as public anyways so I see no real reason to go to another [Place]. I may open up another chat item to keep the conversations separate … There's always this feeling of not wanting to exclude people, particularly in a community that is so close and has a rich culture like we do."

When asked about the situations under which they would use different places, most participants hypothesised that they would use different places if they were also involved in distinctly different groups that did not know each other.

We questioned the small group of people involved in the secondary "Games" place. Its creator said:

"I started using it to discuss things that weren't really relevant to all of [the people in the main place]. So, things like games. A lot of people who are in [the main place] don't talk about it, they don't really care about it, they're not part of that group."

In summary, while CB's Places were originally conceived as a way to have groups create many different locales, sites and means, this did not reflect how they were used. Instead, the study

group saw a single Place as containing all community members and their activities. Thus our study group was too cohesive to make much use of multiple places. For multiple Places to be useful, we suspect there needs to be separate, distinct communities. Instead, we saw that the group used media items within a Place to implicitly create *mini-locales* used by group subsets. That is, people would post things to the Place that only a few people would be interested in, and let them make their own choices as to what was worth viewing.

4.2.6 Provide a way to relate locales to one another.

The relationships between social worlds are important influences on people's activities. These include containment relationships, such as a department containing research groups, membership relationships, such as a researcher being on two project teams, and so on. The system needs to make the relevant relationships visible to the individual. CB's design was intended to support this by having multiple Places in view, and by letting people belong to multiple places.

As mentioned in the previous section, this group used only one primary place. Indeed, people who did create the other places did so because its purpose was quite distinctive from the primary place. Yet even in this limited use of multiple Places, people expressed an interface issue as revealed in this quote:

"In the Game place, there were the same people as in [the main] place so I had the same camera picture twice [for each person] and it was just totally cluttered … so I didn't like it that much."

Another limitation is that Places were hard to bootstrap because people in the primary place did not receive notification of when a new place was created, i.e., new places became 'by explicit invitation only'. Even if they accidentally discovered a new place by its listing on a popup menu, they could not find out anything about it unless they entered it. This was something they were very hesitant to do uninvited, regardless of their curiosity:

"I noticed them but I didn't go into them because I wasn't sure who they were and so I wasn't sure if I was invited. Like I was curious about the games place but it was like, oh well, I don't know who set it up and I don't really know what it's for, so I'm not going to join."

As mentioned, people did create 'mini-locales' within a Place. Yet there was no explicit way to relate these together except by semantic content. For example, a person may create a mini-locale by posting a web page and starting a chat about its contents, but the items that contained these could be scattered around the sidebar.

Thus CB's design did not satisfy this principle. Places appeared to define distinctive communities, with little need to show the relationships between them. While people exploited media items to create mini-locales, they had no way to cluster related items together, e.g., by spatial positioning or grouping mechanisms. This clearly needs to be addressed in future versions.

4.2.7 Allow individual views.

People commonly work in multiple locales simultaneously. Their collection of locales, particular to the individual, is called a *viewset*. Additionally, each person has their own individual perspective on each locale dependent on their role and involvement with that locale, called a *view*. Hence the system needs to provide tailorable viewsets consisting of multiple variable views of locales. CB's design was intended to promote these individual views through Places and through its Focus controls (discussed later).

The few participants who made use of more than one place were able to easily discern the different contexts across their viewset and direct communication appropriately. Apart from the problem of repeated presence tiles described previously, people appeared comfortable with the idea of multiple places.

People were also comfortable with their individual views within a place, i.e., they all knew that item display was entirely local, while the information within them was shared. That is, each person had their own idiosyncratic view of the order of items on the sidebar, the size of items, the individual focus settings, that the raising of tooltip grandes and full views was entirely local, and that the owner of an item could see owner-specific information and controls.

Many participants' comments indicated that they wanted a greater level of control over their individual views than CB currently provides. Some people asked for specific features such as to be able to move media items within the sidebar so they could place the most important information at the top. More generally we received comments like:

"If I could have a place where it has all the individuals where I only want to maintain a slight amount of awareness, and a place where it has a lot of awareness."

This comment reflects not only a stronger need for personal view control, but that part of this desire stems from people wanting to group things into mini-locales of varying interests to them.

4.2.8 Allow people to manage and stay aware of their evolving interactions over time.

Awareness of past actions and outcomes, present situations, and visions for the future are important for creating plans and strategies. In addition, current information is time critical and must be kept up to date. Within CB, only the Chat item was designed to maintain a history of conversations and interaction.

Study participants placed themselves into two distinct groups; one group wanted only information from CB about "right now", while the other group wanted more long term history information.

The "right now" group did not think that time information would help them gain any further awareness, or felt that the extra information that they would get would not be of any use to them.

"There's not much history but the things I use it for I wouldn't necessarily want history … to know if someone's there, to know if they're busy or on the phone."

Yet this reaction was item specific: members of this group still asked for timestamps to be added to Chat messages:

"That would be nice, if it had a time stamp on messages cause I'll see a conversation when I log on … and I don't really know when so I don't know if I can add … like if it was six hours ago I'm not going to jump in but if it was five minutes ago I might"

The other group felt that historical information could help them to predict useful future events:

"If I look to see if [Participant]'s around and see he's not, I have no idea of when he left which could be a good indicator of when he's probably coming back."

They also wanted presence information to be augmented with known future events:

"I can't get any sort of long term prediction because I only get a small snapshot. Like maybe if CB videos were augmented with a calendar of when the person had appointments that day I could know they're available right now, they don't look like their busy, but they have a meeting in 10 minutes."

Another interesting pattern occurs when participants were away from their desk for a while. When they came back they were unable to assess whether they should respond to new information..

"There were times when there was a message in there… like I didn't know when the person had said it so I didn't know if it was relevant to reply."

The long-term awareness group also suggested interest in long-term information about other users' rhythms [2], because this would help them predict future behaviour.

In summary, while people's reactions varied, several uses of an evolving interaction emerged: to see if lingering items are still relevant, to review past activities, and to help form predictions of future activities so that they can manage their interaction more effectively.

4.2.9 Provide methods for controlling focus.

Focus refers to how people direct their attention, determining their awareness [16]. As an individual interacts with different people, places and artefacts over time, their focus of attention is constantly changing. As an awareness system, CB was also designed to also allow individuals to change their focus on the people, places and artefact items in the sidebar. Of course, raising the tooltip grande and full view let people add focus by drilling down; the success of this was discussed previously. Another means was to use a focus control – a slider – to adjust the size, information content, and overall awareness of a tile.

Focus controls were primarily used to increase awareness of video images in the Presence tile, where people would set it at their maximum size. To explain, when there are too many items to fit in the sidebar, CB automatically reduces the size of all existing items to make room for the new one. Yet people wanted video items at full size. One user expressed the frustration like this:

"I don't have the time to go in and adjust people's focus all the time… I have tried in the past but then all of a sudden I'll have too many items on my bar and then CB will re-adjust everybody's focus and then it's like, why did I bother in the first place?"

Another problem is that CB did not increase the size of tiles when room became available:

"You know, lots of people log in and it makes everybody smaller then some of them leave and so people are sort of arbitrarily sized so there were certain people I would go back and make them bigger so that I could actually see them."

"I'm just getting rid of white space so that everyone will fit - I actually wish that CB would do this for me but it doesn't."

Sometimes, however, people used the focus to reduce the size of items that were less interesting, e.g., people's photos down to names, or empty chats. Overall, they used this strategy to make the awareness information on the sidebar more viewable 'at a glance', i.e., a quick glance at it would let them concentrate on the items of most personal interest.

People also used focus to mute items that had already been read such as a web page. *"[I shrunk it] because I read the contents and didn't need to see the whole thing anymore until it changed."* This behaviour suggests that users had a clear preference for seeing constantly changing dynamic information on the sidebar, rather than static information that had already been seen.

Obviously, CB's focus control is flawed. People should use it to adjust awareness needs rather than fight the constraints imposed by automatic resizing. This should be changed. For example, people could indicate that video should always be as large as possible. Automatic resizing should reuse empty space when the bar is less busy. User settings should be stored so the system remembers preferences. Grouping, as mentioned previously, could be used to place and shrink less interesting items.

In summary, focus control in real life is a matter of glancing around and attending to things as desired. This is reasonably emulated in the tooltip grande and full view. However, the explicit focus control in the tile view, while valuable, is too awkward and too much work.

4.2.10 Provide methods for controlling nimbus.

In real life, an individual adjusts how they are visible to others, i.e., positioning oneself to include others in an interaction, or restricting what others can see because of privacy concerns. This adjustment is called *nimbus* [16][15]. CB's Presence item had a nimbus control, where people could limit how others viewed them, i.e., as video, as a photo, as a text title, or by on-line activity.

People with a webcam typically left their Presence nimbus at its maximum, i.e., to show video. Those who did not own a webcam changed the nimbus to a static picture instead of a test pattern (in hindsight, this should have been done automatically by CB).

People used reciprocity as the reason for keeping their nimbus setting at the richest level possible. First, video was perceived as so much richer and more useful than the other presence information (as already discussed in subsection 4.2.4) and users wanted to provide others with good presence information. The second reason is a function of the community's social practices, represented by this one user's comment about reducing his nimbus:

"The social environment was such that it would be weird if you [reduced nimbus]… People may ask questions like why."

One home telecommuter would adjust nimbus when he left the home office, as he was concerned that other family members using the room would be caught on camera. However, he always readjusted the nimbus back to full on his return.

Even when people did not use the nimbus control, they still believed this power was important. Indeed, some people did not join CB because they did not feel comfortable being on camera all the time. Some were concerned over being caught on camera at embarrassing moments. Some telecommuters didn't think that the camera was appropriate for their home environment.

CB's nimbus control was also too discrete: either video was transmitted, or not. People wanted finer control, where video could be transmitted but at reduced fidelity. One of our interviewees described his problem and solution:

"When I'm at the university I basically don't care what people see of me because I'm in a public place … but when I go home

I'm very conscious of what people see of me because I'm not constrained like I am at work, you know to be dressed appropriately [and] there's other people at home that don't like cameras … I don't actually change my nimbus though when I'm at home, it stays at full I'm pretty sure because I still like to present people with a video as oppose to the picture because I don't feel the pictures provide much information, so what I instead do is take my camera and I adjust the focus … so I'm blurred in the background."

We also saw people pointing the camera at the keyboard or mouse, affording privacy while still providing some awareness information to others. Even when neither of these techniques was used, all participants made sure that their camera was directed so as not to capture people just passing by or visitors.

In summary, CB provides explicit but limited control over nimbus. People worked around this by implementing their own fine-grain adjustments of video nimbus by changing their environment rather than using CB's less than perfect controls.

5. DISCUSSION AND CONCLUSIONS

Because there are overlaps in the above principles and study results, we will structure our study implications within four themes [12].

Multiple social worlds through public locales. Principles 3, 5 and 6 (Table 1 and Section 4) argue that each person concurrently 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.

While we know that multiple places are theoretically sound and have proven useful in another casual interaction system [7], CB Places were used very little. As mentioned, a CB Place seemed to define a community rather than a public locale. Instead, people tended to use collections of CB media items as implicit locales; this led to disparities between the levels of awareness people wanted on the items defining these locales.

The Locales Framework [6] strongly features the concepts of *centre* and *peripheries*. The peripheries of our study's social world became very apparent during our interviews. People felt on the periphery for a variety of reasons: not having a webcam when most others did; not being collocated with the majority of the group; and not being part of the "core" group consisting of the Professor and his students. Those on the periphery were unable to express their distance from the centre to others and the system did not support their different needs within the space.

Clearly, CB has to be redesigned. Perhaps these mini-locales can be created in place, with items easily grouped and its nimbus easily adjusted. In this way, those at the centre can have greater awareness and opportunities for interaction. Those at the periphery can still acquire passing information about the core group's activities, and move towards the centre as desired by readjusting their focus (see the last part of this section).

Ad hoc groups. 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). We believe that all of these points indicate that support for ad hoc groups is necessary in a casual interaction system. Ad hoc groups require lightweight facilities – meaning easy, quick and simple – for creating, joining, and populating (both people and content) a locale.

As previously mentioned, CB Places did not encourage ad hoc groups. Even when a new Place was created, it was not used that much and it was not long lived. This could be partly explained through interface issues, e.g., that people were not notified when a place was created, and that they were reluctant to enter a place unless explicitly invited.

However, a more likely interpretation is that ad hoc groups *were* created, but outside the explicit structure offered by the Community Bar. Just as people created their own implicit locales through media items within a place, they also formed implicit ad hoc groups as a function of their awareness and CB activities. This was evident by the way chat items were used. Typically, only subgroups partook in discussions in chat items, and different chat items were often created (or taken over) for different purposes and people. Similarly, different sub-groups were interested in different things at different times: this likely led to some of the differences in how people interpreted some media item awareness information as useful *vs.* as clutter and distracting. People seemed comfortable – even those who were 'on the periphery' – of doing all this ad hoc group formation within the context of the larger CB community.

Implementing the CB changes mentioned previously could also help how people who are part of a larger community form and dissolve ad hoc groups within it. Rather than create an explicit structure for ad hoc groups, the easy creation and use of a mini-locale may be all that is needed to define it.

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). 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.

The design of peripheral awareness with transition to foreground interaction is borrowed from the *quick drill-down and escape* design principle in the Sideshow system [4]. The success of the Sideshow deployment trial indicated that this design principle should work, and this is what we saw in CB. No one mentioned problems with the drill-down awareness-to-interaction-transition design, suggesting it is a natural way to move between the two.

Our usage data and interview results indicate that the primary use of CB is as a peripheral awareness tool that promotes feelings of belonging to the group. Our data also shows that it is the Chat item's full view that is used most often. CB succeeds because awareness can be easily transformed into intensive, full-focus collaborative interaction. Of course, CB could do much better in supporting rich interaction, e.g., by offering high quality Voice over IP, and better groupware interaction tools.

Focus/nimbus control of centre/periphery relationships and awareness. Most groupware systems give people a binary choice for their involvement, i.e., they are either in or out. If out, they do not receive any awareness information (but see [3] for an exception). An example is groupware using a rooms metaphor, where people within a room all have opportunity to interact with one another, but as soon as a person leaves that room, that activity becomes invisible [9]. This binary approach to participant involvement is directly contradicted by our principles. In particular, we believe the centre and periphery relationship – a central tenet of the Locales Framework – can be represented by the focus/nimbus model [16] (Principles 9-10). Awareness 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 how people stay aware of interactions over time.

The study uncovered many centre/periphery issues. Of most concern, it revealed that people on the periphery of the core group were inadequately supported. They were unhappy about the notifications they received, as these were usually about communications more relevant to people near the group centre. Their concerns indicate that CB's focus and nimbus controls are not sufficient for peripheral group members.

The central issue seems to be that, despite our efforts to do away with the rooms metaphor, CB Places are still too much like rooms. We feel strongly, after seeing these results, that CB has to provide more *transitionary zones* if it is to be effective. Leveraging the idea of implicit locales, these zones could be constructed as mini-places that allow some of the interaction within it to "leak" out into the surrounding area. Perhaps the idea of focus and nimbus controls for awareness is the wrong approach. Instead, people need lightweight, natural and flexible control over how they move between the centre and the periphery. Awareness is adjusted as a consequence.

Final words. While we have been somewhat critical of CB throughout this discussion, it is important to remember that CB is fundamentally useful. Our group has voluntarily used CB for over two years. CB's design also evolved in this period in response to the feedback received from dedicated and enthusiastic users. The problems reported here are yet another way for us to reflect on how we can move CB's design forward.

The basic repeating pattern throughout our results was that the theoretical principles are valid. When people could use CB's design features in a way that matched the principles, we received very positive feedback. When CB's design – in spite of good intentions – was a mismatch to this theory, our study participants would point out that we were missing such support, and would often find workarounds to implement this support 'outside' the system's structure.

Acknowledgements. This work was partially funded by the NECTAR NSERC Research Networks Grant. Special thanks to all our study participants, who volunteered their time.

6. 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] Begole, J., Tang, J., Smith, R., Yankelovich, N. Work Rhythms: Analyzing Awareness Histories of Distributed Groups. *Proc. ACM CSCW* (2002), 334-343.

[3] 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.

[4] Cadiz, JJ, Venolia, G.D., Jancke, G., and Gupta, A. Designing and deploying an information awareness interface. *Proc ACM CSCW* (2002), 314-323.

[5] Curtis, P., Nichols, D. A. MUDs Grow Up: Social Virtual Reality in the Real World. *Proc 39th IEEE COMPCON* (1994), 193-200.

[6] Fitzpatrick, G. The Locales Framework: Understanding and Designing for Wicked Problems. Kluwer Academic Publishers, (2003).

[7] Fitzpatrick, G., Parsowith, S., Segall, B., Kaplan, S. Tickertape: awareness in a single line. *Proc ACM CHI* (1998) 281–282.

[8] 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.

[9] 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)

[10] Greenberg, S. and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. *Proc ACM CHI*, (2001), 515-521.

[11] 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.

[12] McEwan, G. and Greenberg, S. Supporting Social Worlds with the Community Bar, *Proc. ACM CSCW* (2005), 21-30.

[13] McEwan, G., Greenberg, S., Rounding, M. and Boyle, M. Groupware Plug-ins: A Case Study of Extending Collaboration Functionality through Media Items. Report 2006-822-15, Dept Comp. Science, U. Calgary, Canada 2006

[14] Nardi, B.A., Whittaker, S., and Bradner, E. Interaction and Outeraction: Instant Messaging in Action, *Proc ACM CSCW* (2000), 79-89.

[15] Reynard, G., Benford, S., Greenhalgh, C. and Heath, C. Awareness Driven Video Quality of Service in Collaborative Virtual Environments, *Proc. ACM CHI,* 1998.

[16] Rodden, T. Populating the Application: A Model of Awareness for Cooperative Applications. *Proc. ACM CHI,* 1996, 88-96.

[17] Rounding, M. *Informal Awareness and Casual Interaction with the Notification Collage*. M.Sc. Thesis, University of Calgary, Alberta, CANADA, 2004.

[18] 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.