

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
HxI Initiative & NICTA
Australian Technology Park
Eveleigh, NSW, Australia 1430
+61-2-8374-5581
gregor.mcewan@nicta.com.au

Saul Greenberg
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2N 1N4
+1-403-220-6087
saul.greenberg@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 a field study of a group's on-going CB use. We use its results to reflect upon the matches and mis-matches that occurred between the theoretical and actual usage behaviors anticipated by our design principles vs. those observed in our deployment. As a critique, this reflection is an important iterative step in recognizing flaws not just as usability problems, but as an incorrect translation of theory into design that can be re-analyzed from a theoretical perspective.

Categories and Subject Descriptors

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

General Terms

Design, Experimentation, 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 [9,14]. 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, in an appropriate manner [9], and in ways that easily take advantage of near-by work artifacts to

progress naturally to artifact-centric work. These interactions 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. However, the same studies also found that casual interactions severely drop off when people are physically separated by even small distances [9,14]. Thus distributed communities of co-workers miss out on these valuable interaction opportunities. In response, developers have designed groupware that displays informal awareness information leading to casual interaction between distributed group members, e.g., Instant Messengers (IM) [12], chat rooms / MUDS [4], and video-based media spaces [1].

These tools prove valuable in practice. For example, while IM provides only a rudimentary indication of other people's on-line presence along with easy invocation of a crude communication channel, this suffices to stimulate many casual chats. 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. IM treats one's social communities as a disparate set of buddy lists, favoring isolated chats between two people. Chat groups and their variants have rigid notions of how groups and their memberships are defined, how people present themselves to others, and how conversations are publicized. From a social science perspective, communities are far more dynamic.

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 behavior [9,14] and social science theory [5,13]. In particular, we are motivated by the Locales Framework [5], one of the few comprehensive theoretical group interaction frameworks in computer science, as well as the Focus and Nimbus model of awareness [13]. We have previously derived and combined tenets from these studies and theories into a set of design principles [10]; these are summarized in our results section but not described further due to lack of space. We then used these principles to design the Community Bar (CB), a 3rd generation groupware tool (predecessors in [7,8]) that supplies groups with rich awareness information leading to casual interaction [10,11]. The question is: Did we succeed in applying theory to practice, and if not, why not?

In this paper, we present a field study of a group's on-going use of CB, and use its results to reflect upon the matches and mismatches that occurred between the theoretical and actual usage behaviors anticipated by our design principles vs. those observed in our deployment. As a critique, this reflection is an

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.

ACM Group '07, November 4-7, 2007, Sanibel Island, Florida, USA.
Copyright 2007 ACM 1-58113-000-0/00/0004...\$5.00.

important iterative step in recognizing flaws not just as usability problems, but as an incorrect translation of theory into design that can be re-analyzed from a theoretical perspective

2. COMMUNITY BAR

Community Bar is fully described in [10,11]. This is just a summary of key interface components. Figure 1 illustrates a screen snapshot of CB. CB presents itself as a sidebar peripheral display [3], divided into *Places*. Each Place represents a sub-group, their communication, 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 is the *site and means* [5] for group awareness and communication. Figure 1 shows an individual's view of four Places: CSCW class, G-place, ilab, and mike test.

Each Place contains a number of *media items* [8], representing people (as live video, photos or names), public conversations (as chat dialogues or sticky notes), or group artefacts (e.g., photos and web pages). The Place names, the membership of people to that place, the choice of media items and the content of these items are completely defined by the group on a moment by moment basis. Media items have three different forms: the *Tile* view, the *Tooltip Grande*, and the *Full* view.

Tiles, meant for peripheral awareness, are always visible in the sidebar. Thus all members within a place will see those tiles. For example, all people in the 'mike test' place 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 shows 5 different Tile types representing people, conversations, or shared information.

When a Tile's contents capture a person's attention, he or she can explore and interact with that information in greater detail. First, when the person mouses over an item in the bar, CB displays a transient *Tooltip Grande* [3]. 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 the *Tooltip Grande*, the *Full* view appears as a new separate window that 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 an audio conversation. Gregor sees this view 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. The chat item's Tile shows the last few messages, its *Tooltip Grande* shows more messages and allows message sending, while its *Full* view shows

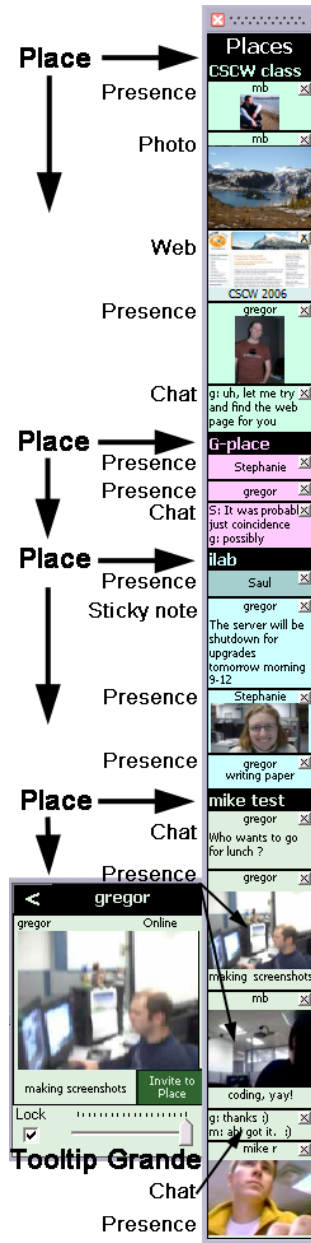


Figure 1. Community Bar

all messages, the place members, and who is 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); this implements and therefore subsumes most capabilities of the Notification Collage (NC) [8].

All *Tooltip Grandes* contain a 'focus' slider control (seen in the *Tooltip Grande* in Figure 1) that allows the user to control their personal view of Tiles [13]. Moving the slider from right to left not only shrinks the Tile's size in the bar, but also semantically changes the information. Similarly, the 'owner' of a presence media item (i.e., the person that created it) can adjust a 'nimbus' slider control in the *Full* view to limit the level of detail that others can see [13]. Thus what a person actually sees in a Tile is constrained by both focus and nimbus.

3. FIELD STUDY

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 activities. 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 cultural norms. We were interested in examining how the group used CB after this period of adoption, i.e., after their social practices had stabilized. This suggests that a naturalistic longitudinal field study was needed.

The logistics of bootstrapping groups with CB and then subsequently monitoring them for many months is onerous. Instead, we adopted a pragmatic approach and studied two groups. Group A was an 'in-house' group that had developed a culture of CB use for several years; while including the two creators of CB, the majority were members of the surrounding research group uninvolved with CB's development. Group B was a smaller outside commercial software development group who had been using CB for only a short time. This choice deserves discussion.

Because Group A is in-house, its members may be biased towards CB. Yet we stress that the group is very worthy of study:

- The group had been using CB for several years, and had thus developed a rich culture around it.
- Group membership had changed over years as people came and went, and thus went beyond the original core group that had vested interests in it, i.e., people used it because they wanted to rather than because they were asked to use it as experimental subjects.
- Most current members worked on quite different projects and were uninvolved in the actual CB research.

- All were experienced with groupware research, and thus better able to reflect on their practices.
- The principle investigator of this field study was *not* part of this group: her involvement was for the express purposes of setting up and conducting the field study.
- At the extreme (i.e., if significant favorable biases did exist), this group represents a ‘best case’ condition; thus any problems seen would likely be reflected in other groups.
- Our opinion is that most members were not biased one way or another towards CB; they saw it merely as another communication tool at their disposal.

Group B is our reality check. While Group B did not have enough members or lengthy history of use to make it a rich case study by itself, it sufficed to see if the general phenomenon we observed with Group A were also present in Group B, which in turn would suggest that any bias effects were not significant.

Participants in Group A comprised 15 people. All had real world work and social relationships with each other. 11 were current members of a research laboratory at a university, while others worked elsewhere. The breakdown was:

- 2 co-creators of CB (a professor and a graduate student),
- 4 other current graduate students supervised by this professor
- 5 graduate students supervised by other professors who did not use CB
- 3 former graduate students of this professor who wanted to remain in touch with this group: 2 worked at private industry firms (one located in another city), while the other was a student at a different distant university.
- 1 researcher at another university in another city; this person had a weak academic relationship with this group.

Thus 14 people knew each other very well, while the 15th had only met part of the group a few times. the 11 current lab members frequently saw each other face to face and often participated in social activities together, though individuals would also often work remotely. We also stress that a culture of use of CB and its Notification Collage predecessor had evolved over 5 years, where other laboratory members who had come and gone influenced the current culture of use of our study group. That is, the current culture had grown over several generations of users. Finally, CB use by this group preceded the study (i.e., the study did not trigger CB use): 10 had been interacting through CB on a regular basis, 4 less regularly, while the last was a new user.

For Group B, CB was made available to a work group comprising 17 people. After 3 weeks (which was when we interviewed them), 10-15 of them reported that they tried CB, and around 5-7 of them used it daily or almost daily.

Method. For Group A, we conducted an in-depth series of interviews with 9 people, and analyzed people’s diary entries and their activity logs of CB usage.

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

Logging. We instrumented CB to log all actions, including what shared information was posted to it (e.g., chat message contents, web pages), and how individuals used their clients (e.g., raising Full views). Data was processed and interpreted to obtain information on how CB was actually used.

Diary. We created a special diary entry media item. 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. Afterwards, 9 participants were interviewed at length about their impressions and experiences with CB. The 9 included the distant graduate student, the researcher at another university, the professor, 4 of his graduate students, and 3 of the other graduate students. The graduate student cocreator of CB was not interviewed. Interview responses were then matched with the logging data and diary item entries to give further insight into how explanations matched actual actions.

For Group B, we conducted interviews with 4 people identified as regular CB users. Because of constraints of the workplace. we did not log their activities or ask them to create diary entries.

4. ANALYSIS: THEORY vs. PRACTICE

Our discussion focuses on the match and mis-match between actual CB use vs. use anticipated by our theory-driven design principles described below [10]. We do this by structuring our analysis around each principle. To set the scene, we first analyze login activity over the study period.

Of special importance is that we did *not* find large differences between Groups A and B’s use of CB, i.e., our main observations of use found in the Group A in-depth study were also seen in Group B. That is, Group B served its purpose as a reality check, where its similarities suggest that Group A’s patterns of use were not significantly biased by the fact that it was an in-house group. Consequently, we devote our discussion below to Group A, as we had richer data on them.

We divided each of the 21 study days into morning, afternoon, evening and night. Each period was further divided into two three-hour blocks: morning began at 6am. We recorded a total of 145 out of 168 potential time blocks, capturing 435 hours of CB activity. Data for 3 time periods was lost due to server problems.

In general, only 10 blocks had no one logged onto CB (mostly at night). On average at least 3 people were logged on. Peak attendance centered on weekdays, between 9am-6pm, with an average of 6 but as many as 10 people online around the middle of the day. Yet login activity was not restricted to work hours: there were concentrations of people in the evening blocks. Daytime group membership differed from the evening group on most days: Figure 2 illustrates one example day with a clear change of members (labeled A–J) between afternoon and evening (21:00 hours), where only one member remained throughout the day. Analysis of chat and diary item responses shows daytime activity tended to be work focused, while evening activity typically had interactions centered on socializing and play.




Figure 2: Membership change. Note 15:00-18:00 period

We now turn to our analysis of CB theory vs. practice, where our discussion is centered on CB's design principles. Each principle restates a fundamental assumption derived from theory. Principles 4.1 through 4.4 are from informal awareness and casual interaction research, e.g. [9,14]. Principles 4.1.5 through 4.7 are from Greenberg et al's restatement of the Locale Framework's principles [5] as groupware heuristics [6], and 4.8 to 4.9 are from Rodden's focus/nimbus model of awareness [13].

Each section first briefly states the theory. Second, it recaps how CB developers translated that theory into specific design features. Third, we use our study results to assess whether people used those features as anticipated by theory. If there is a mismatch, we reflect on why this happened. Finally, we discuss briefly the implications for design of both CB and awareness tools in general.

Most of our analysis is centered around the qualitative data derived from our interviews and diary entries. Excepting for the above, the quantitative logging data results are not reported in this paper, as they tended to illustrate usability aspects of CB vs. our main interest in theory vs. practice.

4.1 Awareness Information Should Be Always Visible at the Periphery

Theory. Informal awareness involves information about people's presence, availability and high-level activity. The information is typically located at the periphery, where people selectively move that information to the foreground of their attention when it becomes of interest. Maintaining awareness should not interfere with one's foreground focus on other tasks. Thus while an informal awareness and casual interaction system should constantly display awareness information, it should do so in a way that attracts attention at only the right times.

Design. CB implements this principle through its always visible sidebar design, inspired by the Sideshow system [3] located at the side of the screen (Figure 1). Awareness information is therefore always present at the periphery. Individual Tiles within the sidebar exude subtle attractors when information changes. For example, chat items change color when the conversation is active. Changes are also reflected by the changed information itself, e.g., video frame updates in the presence item, messages in a chat.

Match and mismatch. There are two relevant questions here. Did the design convey peripheral awareness? Was the awareness useful for the community as predicted by theory? Evidence indicates that both answers are 'yes'. First, most (7 out of 9 interviewees) explicitly said their primary use of CB was for awareness. Second, even though actual communications occurred as short, infrequent bursts, people kept CB on their display for much of the time. As screen space is valuable, it had to be performing some useful function for people to leave it running. Third, we asked participants about this trade-off. Three people with single screens said space was an issue, but all others with dual screens verified that the benefit of the awareness information was worth the screen space 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."

However, we also saw that CB's fixed sidebar location and its single tile size neglected the subtleties of how people acted on peripheral information. For example, one person compared CB to its Notification Collage (NC) predecessor [8] that displayed larger notifications on a full-screen:

"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."

Did users find CB's awareness information distracting, which would happen if it was perceived as a foreground vs. peripheral display? To check this, the diary media item routinely asked "Are you distracted by CB at this moment?" In all cases, respondents said they were not distracted, largely as group activity represented in the CB display only placed modest demands for their attention. Changes within tiles were either routine or infrequent:

"No, as ... there is little direct activity other than the video i.e., I am 'up to date'." (Diary Item response)

"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."

One person also said he wanted to somehow direct his postings to sub-groups to avoid distracting others: *"I wish I could send messages so everyone can access them, but only few are [actually] notified"*

In spite of these positive comments, we also found that the balance between awareness vs. distraction is dependant on personal needs. For example, when a new message is posted to a chat item, the tile subtly notifies people of this change by increasing its color saturation. Interview responses to this feature fell into three categories.

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

Overlooked. *"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."* And more generally: *"If I'm not looking at CB I miss notifications."*

Just right. *"I notice when the text items change and I'll go and see if it's a conversation that I want to participate in. If it's not I'll just go back to what I'm doing."*

Some felt compelled to look at changed items even though they were uninvolved in its conversation – this sometimes led to distraction. Other people ignored changes in Tiles when they thought those changes were not pertinent to them. However, they occasionally felt left out if the conversation turned to something they should have seen. Finally, the most positive people were those who felt directly involved in most conversations.

In summary, CB's sidebar design does meet the basic needs of peripheral awareness display. However, its interface does suffer a mismatch as it does not always achieve a subtle balance between peripheral vs. foreground display. This is due both to its fixed position and size format, and the fact that information salience is highly subjective to the person and their context. The consequences are distraction, over-looked opportunities, and a damping of interactions. Yet these are minor, for all people kept using the system.

Implications for Design. The importance of constant peripheral awareness for collaboration is not a new result by any means. The principle was noted in the early media spaces [1], and peripheral awareness of people's online presence provided by Instant Messengers [12] is likely the major factor in their success. We also believe the lack of constant awareness information to be a large factor in the low success of such technology as TeamRooms [7] and MUDs [4] for ad hoc collaboration. This implies that awareness systems should provide peripheral awareness information with distractions at the 'right time'.

We also see that awareness information is valuable enough for people to sacrifice highly valuable screen real estate. This is a slightly stronger result, though again not entirely new, as it was also demonstrated in the evaluation of the Sideshow system [3] for awareness of personal information.

The theoretical importance of peripheral awareness proved true in practice. However, distraction vs. awareness differs between people, and future designs should strive for a better balance between distraction, awareness, and screen resources.

4.2 Allow Lightweight Transitions from Awareness to Interaction

Theory. In the real world, awareness creates opportunities leading to brief but rich interaction [9,14].

Design. CB's drill-down sidebar [3] is explicitly designed to help people move from awareness into interaction. When awareness information in a Tile captures the person's attention, she can quickly act on it through the Tooltip Grande, and in more detail through its Full view.

Match and mismatch. Participants responded extremely favorably to this transitional affordance of the three CB media item views. For example, all said their choice of the three chat item views reflected their desired interaction. They used the tooltip to reveal more of the conversation and to send quick messages. People used the separate view for extended interaction and to view the entire conversation.

"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 Full view."

"[I use the chat Full view] as an easy way to type, to view conversations & to see if the other person is still typing."

However, not all media item types supported an effective progression of information and interaction across their three views. In particular, people would neglect Tooltip Grandes that offered only slightly more information and controls than the tile. In the presence item, for example, the Tooltip Grande shows a marginally larger and slightly faster updating video, but this adds little over its Tile counterpart (Figure 1). One quote 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"

People were not overly bothered by this; they said that it 'was not a big deal' as it was easy to bypass the Tooltip Grande while on the way to opening the Full view.

In summary, CB's drill down sidebar appears to be a reasonable mechanism for people to transition from awareness to interaction.

However, it is only fully effective if the information content within these views is tailored to reveal a good progression of information and interaction capabilities.

Implications for Design. In general, CB's practice matched the theory. Even so, some types of information were not being presented to full advantage. The current design of CB imposes a 3-step drill down process on media items, but this is useful only if the information itself is differentiated sufficiently to provide value within each progressive view. If no such differentiation is possible, perhaps one of the steps should be omitted. Or perhaps each media item should be better designed to take advantage of all 3 steps. It may even be possible that, in some cases, more steps will be required ([11] for examples).

In the general case, we have shown that the transition from awareness to interaction can be smoothed by having intermediate steps, which also confirms the original premise of the Sideshow interface [3]. In contrast, some other systems have no way to move from awareness to interaction, make use of only a two step process. Of course, new designs should make the transition as fluid as possible.

4.3 Support Groups of Intimate Collaborators

Theory. Research into informal awareness and casual interaction suggests that the people involved are usually known to each other, where they often interact and work with each other [9,14]. We refer to such small groups as *intimate collaborators*, and these are CB's target users rather than (say) communities of strangers.

Design. CB's design supports small groups, and would not scale to large populations. The sidebar interface can only fit a limited number of Tiles. Media items anticipate intimate collaborators: its contents are public to all, and the information within them (e.g. presence videos, conversations) is not normally what one would wish to share with strangers. CB's design concept of public Place also favors several smaller groups of intimate collaborators vs. one larger and less intimate community.

Match and mismatch. From comments received, CB did support small groups, but not as cleanly as we expected. After analyzing our interview comments, we realized that our daytime CB participants clustered into two somewhat overlapping groups: a "core" group of a Professor and the 5 students he was supervising, and a "peripheral" group of the other 9 CB users. Outside of CB, the core group worked together closely; within CB, they often used it for communication. While the peripheral group also interacted frequently outside of CB, they had less work ties to the first group and thus felt less involved with the overall CB community. These different relationships 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. To illustrate, 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 somewhat like outsiders, and that most of the explicit communication on CB did not involve them. As one member said:

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

Yet even these peripheral members realized that their feeling of involvement had more to do with the group makeup than the system implementation:

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

We should stress that participants said they did feel like part of the larger group and that, whether core or peripheral, all expressed sentiments on how useful CB was for maintaining an idea of what the group was doing. That is, even though people may have reported that they were on the periphery, they still felt engaged enough with the others to keep CB on display.

A third more socially oriented group arose as day moved into evening (Figure 2). This new group actually involved a mix of both core and peripheral daytime members from the daytime group, although the way they interacted (social vs. work) differed.

Of course, there is a fourth group: those who did not use CB at all, but were still involved with that community in the real world. A valid question is why these people did not use CB. One study participant who also worked closely with a non-CB group commented that it was a matter of social cohesion, or lack thereof:

"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 we see that CB works best for small coherent groups of intimate collaborators. It works less well for people who are peripheral but still "intimate enough." If people are not part of a cohesive group, they do not see CB as a panacea for bringing it together. Thus while CB does match in how it supports intimate collaborators, it mismatches somewhat in how it deals with those who want different level of engagement with the others.

Implications for Design. The results relating to this principle point to important issues about group structure. We saw multiple groups, each with core and peripheral members, and there were no clear boundaries between them. These results speak strongly to the principle of centres and peripheries from the Locales Framework [5]. We discuss these results in combination with results from other principles in the "Locales" subsection.

4.4 Provide Rich Information Sources and Communication Channels.

Theory. In a collocated physical environment, people use a wide array of rich awareness and interaction channels; they can talk, see, hear, smell and touch and share artifacts in many subtle and varied ways.

Design. CB does not attempt to capture all capabilities of the everyday world. Rather, its design favors rich presence indicators (e.g., low frame rate video of people, activity indicators), multimedia artifact sharing (e.g., photos, web pages), and text-based chats. Yet we stress that CB is not limited to the media items displayed in Figure 1 and used in this study: CB is also a toolkit where 3rd party developers can create their own media

items as plug-ins [11]. However, only the media items shown in Figure 1 were available to study participants.

Match and mismatch. People used the chat channel heavily for communication. Unlike IM, the public nature of these chats meant that different group members would join in on conversations as they happened. People occasionally posted photos and web items, and these too triggered exploration by others (e.g., raise the web page or a larger view of the photo), and further conversation.

However, it is the presence item that stands out. 7 participants reported that the rich awareness provided by the video snapshots was the primary motivator for CB use.

"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 of video was a mixed blessing, as it overshadowed those people who used static pictures (e.g., if they did not have a webcam). Participants reported that they often neglected others using these lesser forms of presence information.

"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. I really like the fact that lots of people do use the video."

Two participants who did not portray themselves through video said they felt somewhat left out. One said that people without cameras were like "2nd class citizens." Another, who only gained a webcam in the middle of the study period, commented on this feeling of inclusion:

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

Yet video presence also caused some frustration. While people would try to contact someone they could see on the video, they could not always attract their attention (see 4.2.2).

"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 rich information and communication channels match expectations of increased awareness and interaction. The mismatch is that it is still a far cry from real world interaction. Media item offerings are limited (but can be expanded [11]), and the opportunities perceived by having that information on display are not necessarily realized by its functionality (i.e., what people want to do does not always match what a media item can do).

Implications for Design. Three themes arise from these results. First, rich information awareness and communication was beneficial, and thus the theory is validated not only for CB, but for group awareness systems in general. Second, presence disparities (video vs. no video) led to some negative effect. IM tools solve this by taking a "lowest common denominator" approach, where everyone provides the same low fidelity

presence information but with no opportunity to go beyond that. Obviously, new system designs should increase presence fidelity while somehow balancing disparities arising from technical barriers. Another important implication for design is that awareness should always be coupled with means for communication.

Third, there is the disparity between perceived opportunities arising from what is visible vs. how people are limited from pursuing these opportunities by an item's actual functionality. The implication is that each media item – and awareness systems in general - must be designed as a whole, where one must consider how people can follow through on the information they see and the communications they want to pursue.

4.5 Provide Centres (Locales)

Theory. Collaboration involves groups of people working together for a common purpose. The Locales Framework [5] calls these groups social worlds. Social worlds use *sites* (locations), and multiple tools, or *means*, to work towards their purpose. The combination of social world with site and means is a *locale*. Each person participates in multiple social worlds with many sites and means, so technology must provide and manage many locales.

Design. CB's Places emulate multiple locales, providing the site and means for the members of a social world.

Match and mismatch. CB's Places did not match our theoretical expectations. All participants primarily used a single place. While three other places were created, they were used infrequently by few people, and were short-lived. On analyzing what happened within the single Place, we observed many instances where we thought another place would have been warranted but was not created, e.g., the three different groups mentioned in 4.2.3, divergence in conversational topics, and instances of subgroups working closely together for an extended period of time.

When asked why they did not create new Places, participants responded in 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 CB, 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."

Further analysis revealed what was really happening.

1. People saw Places as a way to *define and separate communities*. When asked, they said they would use different Places if they were involved in distinctly different groups that did not know each other.
2. People actually created *multiple 'mini-locales' within a single Place* by how they used media items. That is, people would post information and conversations to the Place that only a few people would be interested in, and sub-groups would form temporarily around this 'mini-locale' by how they monitored and interacted over them.

In summary, while CB's Places were originally conceived as a way to have groups create many different locales supporting

various social worlds, this did not match how people conceived them as ways to define and separate communities. Yet multiple locales were created within a single Place, verifying the basic theory. We now realize that Places are too heavy-weight; locales are something that may form, evolve and disappear rapidly within a community, and this may be a consequence of implicit actions. The explicit construction of a Place to hold these Locales proved a barrier to this process. In contrast, the use of media items supported these light-weight locales, but often at the cost of distracting others.

Implications for Design. The community found the public Place of CB valuable, and it served as a locale. However, CB redesigns and other future systems need to be far more lightweight in how they let people create such locales. Our findings suggest that locale creation could perhaps result from people's implicit actions rather than their explicit creation of containers and boundaries as typically found in the room-like metaphors of most public groupware systems [7]. We see also people were comfortable with having multiple locales within a single place, i.e., that perhaps there is no need for hard boundaries between them.

4.6 Relate Locales to One Another

Theory. Relationships between social worlds influence 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 between such groups and locales visible to the individual.

Design. CB's design was intended to support this by letting people belong to multiple places at the same time, and by having these multiple Places in view (as in Figure 1).

Match and mismatch. As mentioned previously, this group mostly used a single primary Place. This limits our analysis, but we can anticipate some problems from the few times people did use multiple Places. One issue is that Places were hard to bootstrap because people in the primary place often did not know when a new place was created unless told about it by another person or by accidentally discovered its name on a popup menu. Even then, they would have to enter it to see what it was about, which was something they were very hesitant to do:

"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."

In contrast, the implicit 'mini-locales' within a Place were always visible, as it was contained within the larger civic structure of the community. People could see what was going on, and they could move from the periphery to the center simply by how they monitored and joined into the activity. Yet CB could do more, as it did not provide and explicit way to group related media items. 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.

In summary, CB's Places feature did not match theory; it isolates locales rather than relate them. However, the mini-locales surrounding media item use did match theory, as it created these relationships. Yet it did so implicitly: relationships were created through semantic meanings and how people used them rather than

through explicit structure. While we believe this *ad hoc* process is a good one, we also recommend that people should have a way to cluster related items together, e.g., by spatial positioning.

Implications for Design. In CB or other related systems, we speculate that we can relate locales to one another by providing more information about them that is visible from outside. That is, locales should be ‘leaky’ where others can get a sense of what they are about. This could include things such as the number and names of members, a short description, or even a sense of the kinds of tools and activities people are engaged in. We know that other communication tools, such as newsgroups and web forums, already implement such features, and that this has proven highly beneficial in attracting new members.

4.7 Allow People to Manage and Stay Aware of Their Evolving Interactions Over Time

Theory. Awareness of past actions and outcomes and how they relate to the present and future are important for creating plans and strategies.

Design. CB’s design for this principle is limited; by default, media items show only the most recent content. Still, the chat item maintains a history of conversations and interaction, and the photo item maintains a record of posted photos. None are time stamped.

Match and mismatch. Participants placed themselves into two distinct groups. One wanted only information from CB about “right now”. The other wanted longer term history information.

The “right now” group said a history would not improve awareness; rather, they make decisions on what they see:

“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.”

The other group believed that historical information could help them to review past activities and 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 thought timestamps would help them discover if chat conversations were ongoing, relevant, or stale:

“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”

This group also believed that integrated past, present, and future knowledge of personal activity [2] would help them predict future behavior. For example,

“Maybe if CB [presence] videos were augmented with a calendar of when the person had appointments that day I could know if they’re available right now. They don’t look like they’re busy, but they have a meeting in 10 minutes.”

In summary, people’s reactions to the value of history varied, thus bringing to question whether this theoretical principle is applicable to all. For those who did believe in it, CB’s primitive history capture was not particularly effective. People not only wanted to review information content, but also when it happened. In essence, they were trying to discover the activity rhythms of

the people around them [2], where they would use this information to manage their interactions more effectively.

Implications for Design. Interaction history can be better incorporated in new designs. In our own newer version of CB, we implemented the simpler features suggested. Time stamps on chat messages was greeted with great enthusiasm by the user community. The larger requests for rhythms and schedules could be implemented as a new CB media item to display communal rhythms, or augmentation of the Presence item to show visualizations similar to those in [2]. The Sideshow system [3] demonstrates integration of calendars into the sidebar interface.

Yet we saw an interesting mix of responses here: everyone wanted some time information such as timestamps on chat messages, while only some people wanted richer time based information about others. We believe that there is a lot of room for further research here concerning the characteristics of the two response groups, what types of time information are always valuable, and how these dynamics change if interactions were asynchronous. Clearly, the way historical information is displayed is important; it should be available to those who want it, but not distracting to those who do not.

4.8 Provide Methods for Controlling Focus

Theory. Focus refers to how people direct their attention, determining their awareness [13]. Focus of attention constantly changes as an individual interacts with different people, places and artifacts over time.

Design. CB allows individuals to change their awareness focus in two ways: first by drill-down into the Tooltip Grande and Full view (discussed previously), and second by increasing and decreasing Tile size / information content through a focus slider control (Figure 1). CB also shrinks tile sizes by automatically decreasing focus if space is at a premium.

Match and mismatch. We saw that focus controls were primarily used to increase, rather than decrease, awareness of video images in the presence tile. 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, but does not increase the size of tiles when room becomes available. Users expressed their frustration:

“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?”

“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. There were certain people I would go back and make bigger so that I could actually see them.”

People sometimes used the focus to reduce the size of less interesting items, or ones already read, e.g., people’s photos down to names, empty chats, or already visited web pages.

“[I reduced the focus] because I read the contents and didn’t need to see the whole thing anymore until it changed.”

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. They had clear preferences for focusing on dynamic information, rather than already seen static information.

In summary, CB's focus control is flawed. We hoped that people would use it to adjust focus subtleties. In practice, they used it mostly to manage space, and to correct imbalances imposed by CBs egalitarian treatment of item sizing, and by its automatic one-way shrinking of items when space became problematic. As well, adjusting focus by moving a slider is too heavy-weight.

Implications for Design. The limited display space available to most awareness systems means that people are often forced to use focus to manage space vs. as a way to increase or decrease an item's salience to reflect personal interest. One solution is to design a dedicated display where space is not an issue. Another is to use global properties, e.g., in CB video presence items should always be as large as possible. User settings should be stored so the system remembers preferences. Perhaps static items whose content has already been seen should fade from view. We should also try to do away with explicit focus controls such as CB's slider, as it is too heavyweight. Ideally, new interaction mechanisms such as attentive interfaces could make focus control resemble the lightweight real life process of glancing around and attending to things as desired.

The design lesson is that focus control has to be a very lightweight, probably implicit, mechanism. This is no easy task. In a group setting, it is very likely that focus is a function of many things, e.g., how willing a person is to be interrupted at a particular moment. In small communities, focus is also affected by others, e.g., when a person believes that some information should be brought to another's attention. Thus an effective focus mechanism could be part of a coordination tool – multiple people could affect an individual's focus, rather than it being a consequence of a single individual's interests.

4.9 Provide methods for controlling nimbus.

Theory. In a physical, collocated environment, an individual adjusts how they are visible to others, i.e., positioning to include others in an interaction, or restricting what others can see for privacy. This representation of self is called nimbus [13].

Design. CB's presence item had a nimbus control, where people adjust a slider to limit how others viewed them, i.e., as video, as a photo, as a text title, or by on-line activity.

Match or mismatch. Most of the time, people set their nimbus on its maximum setting. Those with a webcam typically left their presence nimbus to show video, while those without set the nimbus to a static picture. They said they did this because they wanted to maintain rich reciprocity. First, people perceived video as more useful than the other forms of presence information provided (see 4.2.4). Second, they saw video as a function of the community's social practices, where reducing nimbus could be perceived negatively by others. As one person says:

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

When people did reduce their nimbus, it was usually because of privacy concerns. One home telecommuter adjusted nimbus when he left his bedroom 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 those who did not use the nimbus control believed its provision was important, as it was the only way to regulate the tradeoff of video presence vs. privacy.

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 ... because I still like to present people with a video as oppose to the picture. 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. As well, all participants made sure that their camera was directed in a way that did not accidentally capture people passing by.

In summary, the nimbus control in CB primarily serves as a privacy maintenance mechanism whose use is balanced by expectations and social norms. This does not fully match the theoretical definition of nimbus, but it does reveal that systems such as CB need to consider group concerns over privacy as well as social expectations.

Implications for Design. CB's nimbus control was too coarse and heavyweight, and we saw people work around this by manipulating the actual camera position and its properties. This suggests, perhaps, that systems such as CB should be designed to make the camera the nimbus control vs. a GUI slider. As mentioned earlier, nimbus and focus should be approached as coordination tool for solitude management, with mechanisms that can be easily adapted when the situation demands. Also, focus and nimbus signals should be clearly linked to interaction purposes (e.g., people reported that while they interpreted a blurred video as meaning something about privacy, they wanted more information so they could determine how and when to interact with the person in the scene).

5. DISCUSSION

CB is based on theoretically derived design principles, which in turn has led to a useful – albeit flawed – system. We stress that our critique may over-emphasize the negative; we remind readers that our study group has used CB for over two years and continues to use it. This is very positive, and other groups are now adopting CB in their own ways. The question is not ‘will groupware work’ but ‘can we make it better by understanding its subtleties’? As this study has shown, the devil is in the details.

If a mismatch is seen, where does it come from? Moving from theory to principle to practice can introduce several sources of errors. The theory could be wrong, or the way that theory is encapsulated into a principle could be incorrect. Alternately, the way a designer realizes a principle (and theory) as a screen-based graphical interface may not work as anticipated. This last point is confounded by the fact that the underlying theories are descriptive rather than prescriptive. They do not suggest what designs will work, nor do they truly predict what social behaviors will emerge.

In our field study, we consistently saw the design principles were supported. Where they were implemented well, CB was successful. Where they were implemented poorly, participants often asked for these features, or would develop workarounds to get them. Thus in all mismatches observed between theory and CB's usage, we believe that they arose from how CB's design was realized from the principles. Yet most design flaws cannot be ascribed to simple usability problems. Rather, they point to the deep conceptual difficulties in creating a groupware system that affords the subtleness of social practice.

While our analysis has concerned mismatches arising from CB's design, we strongly believe that similar mismatches can occur in other systems in this genre. Several issues are discussed below.

Locales. One of the larger mismatches that we saw throughout the study was in how people formed and used locales *vs.* how we expected they would be used. Rather than use Places, people instead used small groups of media items to dynamically form, attend, and dissolve mini-locales within a Place. The central issue seems to be that, despite our efforts to do away with the rooms metaphor [7], CB Places are still too much like rooms. This raises questions about other groupware in this genre that tries to 'organize' locales by either spatial metaphors, or by having application-centric gatherings. In essence, we now believe that locale formation has to be an easy, perhaps implicit process that emerges from a community as they do their interactions, rather than from creating an explicit structure *a priori*. Yet designing for such an implicit and tacit process is difficult.

Implicit actions made explicit. Several mismatches occurred between design and theory when we thought that explicit graphical user interface controls could be used to adjust implicit social behaviors. In real life, for example, we adjust our focus and nimbus as a consequence of our actions rather than by changing a setting. Similarly, awareness occurs as a consequence of how we direct our attention, rather than by adjusting controls. Locales are related by how we perceive and move between groups and settings, rather than by explicitly bounding them into some structure. Yet computer systems are poorly equipped to recognize and act upon implicit actions. While we saw that people do manage to mediate what they do, they often do this in spite of the system rather than because of it. Again, this is a design challenge.

Richness. Another aspect that emerges from this study is that the screen-based groupware is a poor realization of real world richness. Yet in spite of this, CB did remarkably well at preserving aspects of group culture. This is good news for groupware developers. Some groupware, such as IM, does well simply because they make very basic social needs possible, *i.e.*, awareness of presence leading to interaction. Thus we can iteratively add richness as the base system already fulfills a need. For example, CB leverages the success of chat rooms by adding awareness, multimedia, and other features. We now see opportunities for iteratively adding even richer capabilities to CB as needed. Examples include a means for increasing awareness by promoting time-based work rhythms of group members [2], better communication channels, better video, and better history.

6. CONCLUSIONS

We learned several lessons from this process of moving between theory, principles, design, evaluation, and reflection. While most

other systems are built from a set of ad hoc design principles, we used the Locales Framework, combined with casual interaction and focus/nimbus theory, to generate what we hoped would be a comprehensive set of high-level principles [10]. While this helped us design CB as a working groupware system in day to day use, its design still has flaws. The combination of theory and field study helped us recognize these flaws as areas where the design did not match theory. Thus we and others can now look back to the theory to consider new designs that may correct our mistakes.

Acknowledgements. Research partially funded by the NECTAR NSERC grant and supported by the HxI Initiative, a national program of research led by CSIRO, DSTO and NICTA. Many thanks to reviewers of earlier version of this paper, whose insightful comments helped us improve it considerably.

7. REFERENCES

- [1] Bly, S., Harrison, S., and Irwin S. Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment, *Comm. ACM*, 3(1), (1993).
- [2] Begole, J., Tang, J., Smith, R., Yankelovich, N. Work Rhythms: Analyzing Awareness Histories of Distributed Groups. *Proc ACM CSCW* (2002), 334-343.
- [3] Cadiz, JJ, Venolia, G., Jancke, G., and Gupta, A. Designing and deploying an information awareness interface. *Proc ACM CSCW* (2002), 314-323.
- [4] Curtis, P. and Nichols, D. 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 (2003).
- [6] Greenberg, S., Fitzpatrick, G., Gutwin, C., and Kaplan, S. Adapting the Locales Framework for Heuristic Evaluation of Groupware. *Australian J Information Systems*, 7, 2 (2000).
- [7] 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, MIT Press. (2003)
- [8] Greenberg, S. and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. *Proc ACM CHI*, (2001), 515-521.
- [9] 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*. LEA Press, 1990.
- [10] McEwan, G. and Greenberg, S. Supporting Social Worlds with the Community Bar, *Proc. ACM Group* (2005), 21-30.
- [11] McEwan, G., Greenberg, S., Rounding, M. and Boyle, M. Groupware Plug-ins: A Case Study of Extending Collaboration Functionality through Media Items. *Proc CollabTech, IPSJ SIG* (2006) 42-47.
- [12] Nardi, B., Whittaker, S., and Bradner, E. Interaction and Outeraction: Instant Messaging in Action, *Proc ACM CSCW* (2000), 79-89.
- [13] Rodden, T. Populating the Application: A Model of Awareness for Cooperative Applications. *Proc ACM CHI* (1996) 88-96.
- [14] 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.