

Common Ground and Small Group Interaction in Large Virtual World Gatherings

N. Sadat Shami, Thomas Erickson, Wendy A. Kellogg
IBM TJ Watson Research Center, United States of America
{sadat, snowfall, wkellogg} @ us.ibm.com

Abstract. Virtual worlds can allow conversational participants to achieve common ground in situations where the information volume and need for clarification is low. We argue in favor of this assertion through an examination of a semi-structured activity among hundreds of users held in a virtual world. Through the idea of implicit grounding, we argue that the affordances of contextualized space, knowledge of the social occasion, and creative self presentation allowed attendees to achieve common ground in a low information volume, low clarification need activity. We use the success of the event to re-examine and extend Clark and Brennan's work on grounding in communication.

Introduction

The term common ground refers to the mutual understanding among communicators about what is being discussed (Clark & Brennan, 1991). The interactive process through which communicators exchange evidence in order to reach mutual understanding is referred to as grounding (Clark & Wilkes-Gibbs, 1986). Establishing common ground has been shown to be essential in the success of collaborative activity (e.g., Kraut et al., 2002).

According to Clark and Brennan's framework of key attributes that constrain grounding in conversation (Clark & Brennan, 1991), virtual worlds should have significant liabilities. For example, even though user avatars are co-present in a virtual world, a user may step away from puppeting her avatar, causing her to 'be

there without really being there.’ Or multi-tasking may achieve the same result. Also, adjacent users in virtual worlds may see different people or objects depending on their camera settings. Technical difficulties aside, these are some of the reasons virtual worlds have not been more successful as a collaboration platform (Erickson et al., 2011).

In keeping with this, we would expect users of virtual world technology to suffer from a lack of common ground and from the extra effort required to achieve it. However, as the results reported in this paper will show, and consistent with the most recent work on grounding constraints in computer mediated communication (CMC), the ability of virtual worlds to support grounding may be less constrained than theory would predict. The affordances of contextualized space, and situational characteristics such as knowledge of the occasion and self-presentation in virtual worlds may facilitate grounding, particularly if the amount and type of information to be communicated and the need for clarification is low.

In this paper, we draw on the concept of ‘grounding needs’ (Birnholtz et al., 2005) – the amount of ambiguity that must be resolved in the negotiation of common ground – and show how grounding was achieved in a semi-structured event held in a virtual world. Understanding the situations and processes that facilitated grounding can assist designers of virtual worlds and virtual events, and improve interactions in virtual worlds.

The rest of the paper is organized as follows. We review related work on common ground in CMC. We then argue how grounding and small group interaction can be achieved in certain situations in virtual worlds and demonstrate this empirically. Finally, we conclude with a discussion and a set of implications.

Related Work

There has been considerable work on how grounding is achieved through various CMC technologies. Research has shown that visual information can reduce the verbal communication needed to achieve grounding by offloading it to non-verbal channels (e.g., Gergle, Kraut & Fussell, 2004; Kraut, Fussell & Siegel, 2003). Veinott and colleagues (1999) found that video helped non-native speakers achieve common ground through non-verbal gestures. Birnholtz and colleagues (2005) show how grounding was achieved through chat in a large, geographically distributed ad hoc group. It seems the only CMC technology left to study in relation to its effect on grounding is virtual world technology.

To the best of our knowledge, there has been little work exploring how the affordances of virtual worlds may help or hinder grounding. Bowers and colleagues (1996) looked at the nature of turn-taking and how user embodiments are used in the MASSIVE system (Greenhalgh & Benford, 1995), but did not investigate the process of grounding. Kelly and colleagues (2004) describe an experiment comparing the process of establishing common ground in shared

visual spaces for real and virtual environments. Traum and Rickel (2002) describe the design of an embodied agent and its dialog model for successfully engaging in conversations, including establishing common ground. Without specifically investigating grounding, several researchers have discussed the issues that arise when virtual worlds are used for collaboration (e.g. Fraser et al., 1999; Heldal et al., 2005; Hindmarsh et al., 1998). We contribute to this literature by identifying situations and processes through which virtual worlds may allow grounding.

There have been various uses of spatial metaphors in structuring interaction among collaborators (see Dourish, 2006 for a review). Examples of such work include media spaces (e.g., Bly, Harrison & Irwin, 1993) and spatial video conferencing (e.g., Sellen, Buxton & Arnott, 1992). Within collaborative virtual environments (CVEs), spatial structures provide a context for reference (Benford et al., 1995) and allow partitioning of the environment into rooms, buildings and zones that may facilitate social interaction (Benford & Greenhalgh, 1997). We draw on these concepts of the use of space and place as a resource for structuring small group interaction in large virtual gatherings.

Extending Clark and Brennan's Common Ground Framework to Virtual Worlds

Grounding in conversation can be achieved through language, as well as physical space (Clark & Marshall, 1981). Linguistic co-presence allows grounding because participants in the conversation are privy to the same utterances. Physical co-presence allows grounding because participants inhabit the same physical setting (Clark & Marshall, 1981). Additionally, participants possess prior knowledge and beliefs, and while those may not be uttered, they nonetheless influence grounding. We will use these ideas of grounding through the physical environment and personal knowledge to develop our argument.

Clark and Brennan present eight properties of media that act as constraints on the grounding process: copresence, visibility, audibility, cotemporality, simultaneity, sequentiality, reviewability, and revisability (Clark & Brennan, 1991). They consider common ground to have been established when "The contributor and his or her partners mutually believe that the partners have understood what the contributor meant to a criterion sufficient for current purposes" (Clark & Brennan, 1991). Birnholtz and colleagues correctly identify that the key phrase here is "sufficient for current purposes" (Birnholtz et al., 2005). They develop the argument that situational characteristics (e.g., asymmetries in access to information by conversational participants), not just media properties, affect *grounding needs*. In particular, they state that the amount and complexity of information needed to be exchanged, and the amount of clarification necessary, define appropriate communication strategies. In situations

where the information to be exchanged is low or uncomplicated and the need for clarification is low, the grounding needs are also low.

We argue that the affordances of contextualized ‘places in spaces’ combined with participants’ understanding of the interaction occasion, and their creative self-presentation in the virtual world provide implicit grounding. We use implicit grounding to refer to the situation where at the *beginning* of a conversation, interlocutors draw on their physical environment, the social situation at hand, and the affordances of the virtual world to create common ground. Participants obviously need to update their common ground moment by moment through questions or clarifications as conversations progress. Nonetheless, grounding can be achieved even when interlocutors do not know each other.

According to the principle of ‘least collaborative effort’ (Clark & Wilkes-Gibbs, 1986), speakers and listeners will strive to use the least amount of joint effort required to achieve their conversational goals. Thus when location in a particular place or participation in a particular event provides evidence of others’ understanding, both speakers and listeners will make use of this evidence to the extent possible to reduce their collaborative effort. For example, if a speaker can see that addressee(s) have gathered in a place reserved for a particular purpose, such as a poster booth, they may rely on that visual evidence instead of producing spoken evidence. If the goal of the interaction is to create new connections and renew existing ones in addition to information exchange, interlocutors will try to make each other feel welcome, to be helpful, and to show interest in each other’s work. Therefore, the use of place and situation not only makes language more efficient, but it may also eliminate the need for some language or actions.

Implicit Grounding in Virtual Worlds

Clark and Brennan (1991) mention two types of coordination required for collective action: content and process. Coordination of content depends on a shared understanding of the subject (know that). Coordination of process depends on a shared understanding of the rules, procedures, timing, and manner in which the interaction will be conducted (know how). Convertino and colleagues (2009) studied content common ground and process common ground in the context of teamwork. Here we use it as a lens to understand interactions in virtual worlds.

Implicit grounding in virtual worlds relies on both content and process common ground. An important type of content common ground is identity information which can be conveyed through a user’s name and avatar appearance. Identity can also include position, background, role and expertise. In social situations where newcomers are involved, identity information is particularly crucial in knowing with whom to connect and in creating a positive impression. For example, the use of real names conveys identity, and can be used for looking someone up in a corporate directory, providing an implicit source of content

common ground. Process common ground can be created through situational information available in a virtual world environment. For instance, in a poster session, it is understood that every poster has a *presenter* whose role is to stand in the vicinity of a poster display and wait for *visitors*. It is understood that a poster's presenter *gives a brief explanation* of the poster to visitors who signal interest by moving into the vicinity of the poster, and that visitors may *interrupt with questions*. It is understood that visitors may *arrive* at any point, and *depart* at any point, and that they will generally move through the posters area looking at different posters. This is important because it enables participants to enter the situation knowing how to act, rather than having to figure it out as the situation unfolds. In such a manner, process common ground enables more efficient communication by drawing on people's pre-existing notions of interaction.

Case Study

Through the case study presented here, we will show how a semi-structured activity requiring low information / low clarification – a poster session – provided conversational grounding for a large number of users and interactions.

Setting

This research was part of a larger evaluation of a distributed conference that used web conferencing and virtual world technologies. The conference brought together the technical leadership of IBM at the annual meeting of an organization known as the Academy of Technology. All 800 members of the Academy, and some guests, were invited to the conference; 502 attended. The conference consisted of plenary sessions conducted through a web conferencing tool; keynotes, poster sessions and socials in Second Life^{®1}; and online text-based discussions. For this study, we focus on the use of Second Life for the poster sessions, as that was considered to be the most successful activity of the conference (Erickson et al., 2011). Our goal is to unpack how participants of a semi-structured activity such as a poster session achieved grounding.

Poster Sessions in Second Life

The Second Life virtual world we studied differed from the public version in two important ways. First, it was accessible only from behind IBM's firewall, limiting participation to authorized employees. Second, avatars were identified with real names taken from the corporate directory.

¹ Second Life is a Registered Trademark of Linden Lab, Inc.



Figure 1. Birds' eye view of the poster halls with the directory in the middle and individual posters along boulevards.



Figure 2. The poster directory displaying thumbnails of posters.

The poster sessions were set along an island on Second Life. Figures 1 and 2 show the setup for poster sessions. As can be seen, the poster area had a central directory with a wall of abstracts that could teleport users to particular posters.

Users could also walk along a boulevard and browse posters at will. This promoted social and unplanned interaction through the introduction of a public and socially neutral area through which people could move en route to their destinations. Poster booths were voice isolated, meaning that conversation in a poster booth did not carry outside the booth. This enabled focused interaction between the poster presenter and attendees. Users walking the boulevard could see who was present in a booth, but could not hear them. This is shown in Figure 3. The boulevard design supported the kinds of conversation that occur in hallways, coffee rooms and lounges in the physical world without interfering with the poster presentations. In total, there were 9 poster sessions held across 11 time zones throughout the 3 days of the conference. Poster presenters were required to present twice; once in their own time zone and once in another. The vast majority of presenters were either newly elected members of the Academy or invited guests. A total of 103 attendees presented posters.



Figure 3. Navigating a poster session. Users could see where activity is occurring (but not hear) and decide to join or not.

Method

We used participant observation and semi-structured interviews. All authors observed the poster sessions, taking field notes and screenshots. We interviewed 30 randomly selected informants using a semi-structured protocol. Interviewees were from 6 countries; 24 were male, 6 female. Interviews lasted 30 to 45 minutes, and all were recorded and transcribed. Data were analyzed by repeatedly listening to recordings and working over transcripts to extract emergent themes.

Results

We use the CoFIRE model (Erickson et al., 2011) to analyze how implicit grounding was created in the poster session. CoFIRE (Coalescence, Focused Interaction, Remixing) begins with Goffman's observations of "accessible encounters" (Goffman, 1963). Here Goffman considers the case of a very large gathering that lacks a common focus (such as a party) within which smaller focused interactions occur (e.g., groups of people chatting). Whereas Goffman is primarily concerned with individual groups, and problems such as how groups maintain focused interaction via processes such as shielding, CoFIRE takes a macro view of the gathering as a whole.

CoFIRE views large gatherings as consisting of three processes that operate in parallel: the coalescence of small groups from the larger gathering; focused interaction among members of a small group; and remixing, where one or more members of a small group disengage and 'browse' the larger gathering, seeking a new group to join. All three of these processes are important because in conjunction they have the potential to expose one person to a steady stream of people. As one attendee said "*Posters are better than a conference call or even a video conference because you can get more people in*" [P22, Male, USA].

Coalescence

One interaction problem that large unfocused gatherings must resolve is how to support the coalescence of groups that are small enough to engage in a focused interaction. The concept of a poster session is, in essence, a solution to this problem. A poster session provides a series of spaces and nuclei for small group interactions, and a set of roles and rules for conducting such interactions. The designers of the posters area went to great lengths to evoke real world poster sessions, and the expectations that accompany them. In the context of the Academy poster sessions, the main interaction problem individuals needed to solve was deciding which poster to attend. While the poster directory – a wall with a series of poster thumbnails (Figure 2) and teleportation provided one

approach, many participants choose to stroll along the boulevard between posters (Figure 1). Some paid attention to where the ‘crowds’ were:

“I could actually see a lot of people hanging out in one poster area which again makes you curious... so you also go to spend time there.” [P15, Male, India]

Others reported being drawn in when they saw people they knew:

“If I saw a person I knew who was looking at a poster that was a natural magnet ... because I knew who this person was and what their interests were.” [P12, Male, USA]

A few said they went in when they saw presenters with no one to talk to because they felt sorry for them. Some participants took extra measures to project their real identity by attaching photos to their avatars or name bubbles (Figure 4).



Figure 4. User appropriation of avatar faces. Users replaced avatar faces with real pictures (left) or let them float on top of their heads (right).

One participant even designed “body armor” out of his slides so that he could present his poster anywhere:

“I attached a large transparent box to my body and put [in] some illustrative slides. Afterwards I bumped into [person’s name]. And I whipped [out] my presentation body armor and showed her some of the slides.” [P17, Male, USA]

Focused Interaction

A second interaction problem that must be solved is how a small group can carry out a focused interaction in the context of a larger gathering. One approach that Goffman refers to is shielding – a way to protect an interaction from the relative chaos of the larger interaction in which it is embedded. In the poster sessions this was achieved by the creation of ‘sonic shields’ around each poster booth so that visitors outside could see what was going on but not hear the discussion. Respondents remarked that this interaction worked really well:

“And I loved that when you left you couldn’t hear, and when you came in you could hear. That was fabulously done.” [P2, Male, USA]

This feature was particularly appreciated because the social gatherings – another activity in the same conference – lacked such shielding (because they

lacked demarcated spaces for small groups) and had too much overlapping talk to allow effective small group conversation (Erickson et al., 2011).

Once one or more visitors were present, the shared understanding of how to conduct a poster session dictated what happened. Presenters took the lead, and often relied on tactics from their real world experience:

“As they came in I’d say hello and ask how they’re doing. So I used the same social techniques that I’d use in the real world to draw people in. ... The only limitation is that you can’t see their expression or read their body language, so that was a little harder.” [P25, Female, USA]

Even so, presenters did what they could to interact ‘naturally’ with their visitors. This is evident from the most common complaint of poster presenters, who reported being annoyed by the awkwardness of having to turn their avatar one way to look at slides and another way to look at the visitors. In face to face situations this is simple to do, but in the virtual environment with its narrower field of view and more cumbersome avatar control, it is not so easy.

One presentation tactic that did work reasonably well was trying to understand to whom a presenter was speaking. This was facilitated by the fact that the Academy conference was designed to display visitors’ real names, and many informants reported looking up the names of their visitors in the corporate directory while they were interacting in the virtual world. As one said:

“I present my poster differently depending on who I am talking to – an engineer I might use more technical information; if I know they’re a top executive in sales I’ll talk very differently to them about how this technical idea might affect their sales. So I felt it was better – I wanted to focus more on the names, and I was writing down names, and matching them with [the corporate directory]” [P25, Female, USA]

To facilitate this sort of behavior, some attendees even wore a floating badge that linked to an online profile.

Remixing

The final interaction problem that must be solved is how a member of a focused interaction can first break off the focused engagement, and second ‘browse’ the gathering to find a new group to join. While the concept of a poster session eases many interaction problems associated with large gatherings, it does little to provide support for breaking off an engagement, other than the shared understanding that interactions around posters are generally brief ones. While visitors might take advantage of the arrival of new visitors to depart, the poster sessions often weren’t that crowded. In the absence of new arrivals, often departure would be mediated by a closing comment like:

“Okay that was nice (presenter name)! So I will leave you to it. Thanks for that and I’ll definitely get in touch with you.”

After leaving, visitors were back in the aisles between posters. Many reported that they moved along the aisles, browsing the posters as they went:

“I walked down the hall and I looked at what the poster topic was and who the presenter was.”
[P9, Female, USA]

This was also a chance to have serendipitous encounters with other visitors:

“I really liked that – I ran into people, and that’s the one thing that makes the virtual world stand out to me. ... You’re not asking them to make time for you, it’s not like scheduling a call... They were there and you were there at the same time: it’s accidental, a serendipity kind of thing.” [P9, Female, USA]

Respondents reported that bumping into people one knew happened most during the poster sessions. It was an informal way of getting to meet colleagues from different countries across the world. Interacting with strangers also occurred the most during these sessions. The informal setting of the poster session and the fact that people were there to mix and talk facilitated conversation. According to one interviewee:

“What makes it different is if there are people there I don’t know, if it’s not a poster session I’m probably not as apt to reach out and say ‘hey, I’m here, talk to me’. I have to have a reason for talking to someone, right?” [P11, Female, USA]

Discussion and Conclusion

In the wake of a study of a large virtual conference (Erickson et al., 2011) we were struck by the participants’ nearly unanimous appreciation of the conference’s poster sessions. Many informants expressed surprise at how well the posters worked, and how similar it felt to a face to face poster session. One said:

“I seriously almost felt like I was there – it was amazing! I’ve done a lot of posters in real life, and you almost got that same feeling because you could see people walking by. When I first got there, someone I knew flew down to say hello and it was just so cool bumping into them. And then when I gave my poster, I could see people walking by, and you got the same feeling: are they gonna stop to see it, are they going to come in?” [P25, Female, USA]

Yet this comment, and others like it, were made despite widespread infelicities ranging from technical problems to the cumbersome nature of gesturing and moving about. Why did posters work so well, in spite of many obvious problems?

This question led to the current study, in which we examined the poster sessions of the conference using the theoretical construct of grounding developed by Clark and his colleagues. In particular, we drew on work by Birnholtz and colleagues (2005) that develops the notion of grounding needs, and argues that in certain situations – when the need for information and clarification is low – quite restricted communications channels can suffice to achieve grounding. This suggested looking at poster sessions with particular attention to their grounding needs.

The paper argues that the success of the poster session is due to implicit grounding – common ground that is established before conversation actually begins. Implicit grounding is achieved in different ways for different types of information. Content common ground was achieved by personal information such

as participants' avatars, and by the ability of other participants' to use the names displayed by the avatars as an index into an associated corporate directory. This information was critical because a basic aim of the conference was to create connections between participants. Content common ground also existed due to Academy members' understanding of the structure and purpose of the annual meeting – for example, the tradition of newly elected members and invited guests giving poster presentations. Process common ground was achieved through the design and collective use of the environment. The virtual environments mimicry of features from face to face poster sessions (slides, booths, aisles), and their use by participants (formally dressed presenters, crowds of visitors in booths, visitors wandering the aisles) helped invoke appropriate behavioral norms. Both types of implicit grounding, by making it clear who was present and what was going on, reduced the need for verbal communication.

An important lesson of this work is that virtual world design is more than window dressing or eye candy. Done properly, it can provide implicit grounding, and thus provide some compensation for the technical and interactional problems of today's virtual world technology.

However, to achieve this and provide implicit grounding, the design must be done strategically. First, designers must be clear on the fundamental goals of the event they are designing. Is the aim to create persistent connections between participants? Is it to enable one to present information to others? Is it meant to foster equal exchange between different participants? Is it meant to entertain? The basic goals, whatever they are, will shape the design. Second, designers must understand their audience, that is, the likely participants in the event and the expectations they may bring to it. The key question is what knowledge are participants likely to have in common? It is this pre-existing mutual knowledge that has the potential to act as common ground. In this case study, poster sessions worked well not because there is something magical about poster sessions, but because participants already knew what poster sessions were and how to behave in them. Finally, designers must create artifacts and environments that evoke their audience's common knowledge and that draw their participants into recognizable patterns of interaction.

References

- Benford, S., Bowers, J., Fahl, L. E., Greenhalgh, C., and Snowdon, D. (1995): '*User embodiment in collaborative virtual environments*', in *Proc. CHI 1995*, ACM Press, pp. 242-249.
- Benford, S., and Greenhalgh, C. (1997): '*Introducing Third Party Objects into the Spatial Model of Interaction*', in *Proc. ECSCW 1997*, Kluwer Academic Publishers, pp. 189-204.
- Birnholtz, J. P., Finholt, T. A., Horn, D. B., and Bae, S. J. (2005): '*Grounding needs: achieving common ground via lightweight chat in large, distributed, ad-hoc groups*', in *Proc. CHI 2005*, ACM Press, pp. 21-30.

- Bly, S. A., Harrison, S. R., and Irwin, S. (1993): 'Media spaces: bringing people together in a video, audio, and computing environment', *Commun. ACM*, vol. 36, no. 1, pp. 28-46.
- Bowers, J., Pycock, J., and O'Brien, J. (1996): 'Talk and embodiment in collaborative virtual environments', in *Proc. CHI 1996*, ACM Press, pp. 58-65.
- Clark, H. H., and Brennan, S. E. (1991): 'Grounding in communication', in L. B. Resnick, J. M. Levine & S. Teasley (eds.): *Perspectives on socially shared cognition* (APA, Washington, 1991, pp. 127-149.
- Clark, H. H., and Marshall, C. E. (1981): 'Definite reference and mutual knowledge.', in A. K. Joshi, B. L. Webber & I. A. Sag (eds.): *Elements of discourse understanding* (Cambridge University Press, Cambridge, 1981, pp. 10-63.
- Clark, H. H., and Wilkes-Gibbs, D. (1986): 'Referring as a collaborative process', *Cognition*, vol. 22, pp. 1-39.
- Convertino, G., Mentis, H. M., Rosson, M. B., Slavkovic, A., and Carroll, J. M. (2009): 'Supporting content and process common ground in computer-supported teamwork', in *Proc. CHI 2009*, ACM Press, pp. 2339-2348.
- Dourish, P. (2006): 'Re-space-ing place: "place" and "space" ten years on', in *Proc. CSCW 2006*, ACM Press, pp. 299-308.
- Erickson, T., Shami, N. S., Kellogg, W. A., and Levine, D. (2011): 'Synchronous Interaction Among Hundreds: An Evaluation of a Conference in an Avatar-based Virtual Environment', in *Proc. CHI 2011*, ACM Press.
- Fraser, M., Benford, S., Hindmarsh, J., and Heath, C. (1999): 'Supporting awareness and interaction through collaborative virtual interfaces', in *Proc. UIST 1999*, ACM Press, pp. 27-36.
- Gergle, D., Kraut, R. E., and Fussell, S. R. (2004): 'Action as language in a shared visual space', in *Proc. CSCW 2004*, ACM Press, pp. 487-496.
- Greenhalgh, C., and Benford, S. (1995): 'MASSIVE: a collaborative virtual environment for teleconferencing', *ACM Trans. Comput.-Hum. Interact.*, vol. 2, no. 3, pp. 239-261.
- Heldal, I., Steed, A., Spante, M., Schroeder, R., Bengtsson, S., and Partanen, M. (2005): 'Successes and Failures in Co-Present Situations', *Presence: Teleoperators & Virtual Environments*, vol. 14, no. 5, pp. 563-579.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S., and Greenhalgh, C. (1998): 'Fragmented interaction: establishing mutual orientation in virtual environments', in *Proc. CSCW 1998*, ACM Press, pp. 217-226.
- Kelly, J. W., Beall, A. C., and Loomis, J. M. (2004): 'Perception of Shared Visual Space: Establishing Common Ground in Real and Virtual Environments', *Presence: Teleoperators & Virtual Environments*, vol. 13, no. 4, pp. 442-450.
- Kraut, R., Fussell, S. R., Brennan, S. E., and Siegel, J. (2002): 'Understanding the effects of proximity on collaboration: Implications for technologies to support remote collaborative work', in P. Hinds & S. Kiesler (eds.): *Distributed Work* (MIT Press, Cambridge, MA, 2002, pp. 137-162.
- Kraut, R., Fussell, S. R., and Siegel, J. (2003): 'Visual information as a conversational resource in collaborative physical tasks', *Hum.-Comput. Interact.*, vol. 18, pp. 13-49.
- Sellen, A., Buxton, B., and Arnott, J. (1992): 'Using spatial cues to improve videoconferencing', in *Proc. CHI 1992*, ACM Press, pp. 651-652.
- Traum, D., and Rickel, J. (2002): 'Embodied agents for multi-party dialogue in immersive virtual worlds', in *Proc. AAMAS 2002*, ACM Press, pp. 766-773.
- Veinott, E. S., Olson, J., Olson, G. M., and Fu, X. (1999): 'Video helps remote work: speakers who need to negotiate common ground benefit from seeing each other', in *Proc. CHI 1999*, ACM Press, pp. 302-309.