Enhancing Collaborative Environments on the Basis of Trust

Stephen Marsh
National Research Council of Canada
steve.marsh@nrc-cnrc.gc**.ca**

Pamela Briggs
University of Northumbria, Newcastle
p.briggs@unn.ac.uk

Waleed Wagealla
University of Strathclyde, Glasgow
waleed.wagealla@cis.strath.ac.uk

Abstract
We conjecture that trust is an invaluable mechanism in the decision making process in situations of risk and uncertainty, and more specifically in how information sharing can be affected, and can affect, trust within groups.

This workshop will examine the role that trust can play in the use, sharing, disclosure and protection of all kinds of information within groups of distributed people and technologies. Within the workshop we will explore questions pertinent to information sharing and trust, automatic sharing of personal (or other) information, privacy, and collaboration.

TRUST AS AN INFORMATION SHARING TOOL

The Challenge: Anticipating the Exploitation and use of Trust by and for Cooperating Users

1. Exploring Trust, A Caveat

The primary goal of this workshop is about how trust can help people work and share together in a collaborative environment, specifically from the point of view of how information as a currency can be linked with trust to facilitate such collaboration. As such, it isn't *about* trust per se.

It has been observed from the literature on trust that most papers about trust, short or long, start with a definition of the concept. At the very least they contain some kind of a definition at some point within them. As a break from tradition, we have decided that this will not be the case here. One reason relates to the sheer number of definitions that exist for trust [Marsh & Dibben, 2002], resulting in a situation where little can be achieved by (re)defining the concept, and adapting another definition is doomed to miss points of view from others that are equally valid. A definition here would muddy the waters – what is needed is an understanding by the reader and participants in the workshop that trust exists in some form (which is in itself culturally dependent) and that it can be used in ways which we shall seek to discover within information sharing in groups. Knowing what it *is* (a daunting, potentially impossible task) is less necessary than knowing that it *exists* and can be seen through its effects. We hope that the workshop participants will bring their own definitions and understandings since, ultimately, it is these that will inform their own work.

2. Trust in a Distributed Setting: Uses and Potentials

Establishing trust within groups of people, even via distributed electronic media, is a relatively complex theme. There exist many studies examining, for example, how face to face meetings can help build initial trust, the role low and high(er) bandwidth communications can play, the ability to judge others across distance-based communications, and so on (for a good overview of these and other concerns, see for example [Zolin et al., 2002], as well as [Zheng et al, 2002, Olson & Olson, 2002, Bos, 2002]).

What is less well examined is the role trust can have within and between these teams in how information is actually *shared*. The disclosure of information, how it propagates through networks of people and machines, and how trust can play a valuable part in, amongst other things, what is disclosed, to whom, when, and for what duration, is what we would like to focus on in this workshop.

3. Discussion

A successful CHI 2004 workshop [Marsh et al., 2004] examined the ways in which trust can play a role in distributed technologies, such as handheld devices, embedded tools, and the Ambient Intelligence that could be

built into these tools. One of the interesting topics to arise from this workshop was the paradox of disclosure. That is, that trust is usually necessary for disclosure of information, but that in general disclosure can help build trust that may not have been there in the first place. This fascinating conundrum, we believe, applies equally to humans and the tools they use, particularly when these tools are themselves distributed and potentially socially aware. It is worth noting here that, in order for trust to be necessary at all, risk would seem to be a necessary precursor (cf [Luhmann, 1979]). This begs the question of what kind of risk, and why.

To a great extent, this question brushes on the concept of privacy also: the sharing of personal information, for example by the devices we use to manage our time, phone books, notes, and so on, is both possible and in some cases sensible. Indeed, a balance needs to be struck between how much collaboration is desired, and how much information should be shared. Deciding when and how much is sensible could be left to the tools themselves with little or not input from their owners, but is this safe? How much of our private information do we want to have shared with others? More interestingly, how much of our private information would we want others to pass on to third parties, and in what circumstances? How can trust be utilized as a measure for the amount of information we are willing to reveal?

While these questions and others are relevant and interesting on a general social front, they attain a much greater significance in the potentially more closely connected workgroup scenario. Decisions made based on trust can potentially affect access to sensitive information, access to tools and capabilities, and access to other people who can help get a specific task accomplished. This is particularly true in distributed settings, while it can apply in more closely-knit groups also. For example, allowing groupware systems to, for example, make decisions about who is capable and should be trusted to perform specific sub-tasks is both non-trivial and could involve the necessary sharing of information with both those people and others in the decision making process (cf. [Jones & Marsh, 1997]).

4. Questions and Concerns

This workshop will address several issues relevant to the concept of trust and the disclosure of and management of information, including:

- What kind of information is necessary to disclose in order to establish trust?
- How much information is 'enough' to establish 'just enough' trust?
- What is 'just enough?'
- What protocols can be defined and used for sharing in different contexts.
- The role of trust and information in establishing and maintaining successful group work.
- What to share, with whom, when, and for what reasons?
- What is the value of tools that can explain why information was shared within the context of trust, given different situational needs.
- Can trust be used to build trust? How is trust transitive in smaller and larger groups?
- How does trust evolve over time? What effect does this have on information disclosure, sharing and use?
- What models of trustable information sharing exist, and how can they cross fertilize each other for maximum impact?
- How does geographical distance affect disclosure of information? What of cultural distance?
- How does distance affect what kind of information is shared to establish trust?

5. Concluding Remarks

The use of trust in group formation and maintenance is not in itself new. However, as a relatively poorly understood concept, trust remains an enigma. At best, we understand trust as we use it; as individuals in a social world. Using trust as a mechanism for decision making is a step beyond this, but one that appears to make sense, particularly as regards the use of information and its disclosure in groups and societies, and how this can ultimately come back to affect trust once more.

As a research topic, trust alone is both challenging and exciting. Coupling this with the practical necessities of its use and application results in a workshop topic that we believe will result in stimulating and exciting dialogue. It is our hope to spur on much more closely knit research in the area as a result of the workshop.

6. Format of the Workshop

6.1 Participation Solicitation and Selection

We aim to bring together practitioners from various fields, such as CSCW in general, autonomous agents, mobile technology, Ambient Intelligence, HCI, and AI. Our goal: a diverse group of practitioners to promote lively discussions. Solicit papers from known practitioners. Solicit participation from book contributors.

Participants will be selected from solicited (via email, web sites, mailing lists, etc.) submissions of up to 10 pages for inclusion in the workshop proceedings (to be distributed at the workshop and available online). Submissions will be peer-reviewed by at least two qualified reviewers.

The maximum number of participants for a successfully interactive workshops is 20. An optimum would be between 15 and 20.

6.2 Method of Interaction – 1 Day Workshop

Games, paper presentations, group discussion on each paper to identify key points, breakout sessions.

Reports from breakout groups and group discussion. A final group session to discuss thoughts and findings, prepare for CSCW'04 poster and/or presentation, decide on schedule for publication in Interactions report. We will present our results as a poster and a SigGROUP bulletin report. This is also the second step to an edited book (the first step, book proposal, is in process).

We propose to finish the day with a dinner for attendees to encourage further collaborations.

6.3 Proposed Workshop Schedule

9.am start. Introduction by organisers (15 mins)

9.15-10.00 (Mis)Trust game – a worked model based on a scenario where each person in the room offers potentially trustworthy and trustable information, to examine information flow, trust levels, and methods of exploitation of information and trust.

10.00-11.00 Brief presentations from the workshop members of their own work, ideas, and thoughts.

11.00-11:15 Coffee break.

11:15-12.30 – brainstorm on directions, definitions, requirements for real trust and understanding in the group.

12.30-13.30 – lunch

13.30-14.00 Summarise brainstorm, identify main questions arising from this and the presentations, as well as extant literature on trust. Divide into (3-4) focus groups to discuss identified individual concepts.

14.30-15.30 focus group discussion.

15.30-15.45 (Working) coffee break

15.45-17.00 Focus groups present discussions. Devise workshop summary, compose poster and report, decide on dissemination methods

17.00-17.30 wrap up and leave.

In the evening we propose a dinner and get together to continue discussion and forge alliances in the field.

6.4 Pre-Workshop Activities

A website with annotated bibliography is being prepared on the topic of trust and discolosure in CSCW. This will be made available to the accepted workshop members in the first instance, and built up by these members (following the workshop, the website will be made public). The website will also contain submissions, discussion pieces, and links to sites of interest to the community.

The design of the Disclosure game to be played in the workshop will be refined from a previous iteration, and any requirements passed to workshop members prior to CSCW.

6.5 Plan for Dissemination

We plan to submit a preliminary report for consideration to SigGROUP Bulletin. We are in the process of negotiating with a major publisher to produce a state of the art volume on the topic, and individual contributors to the workshop will be invited to contribute chapters based on their work.

Technical Requirements

OHP and projector screen required. Our own LCD projector can be supplied.

7. Organisers' Backgrounds

Stephen Marsh is a research scientist at the National Research Council of Canada's Institute for Information Technology, in the Privacy, Security and Trust Team. His current research interests include the application of social-psychological theories to agent architecture design and the development and philosophy of trust and socially adept technologies. He has published several papers in the fields of autonomous agents, HCI, and CSCW, including his PhD thesis containing the first computationally tractable model of trust [Marsh, 1994], and preliminary call to arms for Social Adeptness in Autonomous Agents [Marsh, 1995], and has developed architectures for adaptive information delivery and retrieval, and adaptive eCommerce web sites based on aspects of inter-agent trust and personalisation. He has participated in several workshops, including the CHI 2000 Development Consortium. He has co-organised two CHI workshops, in 2002 on Socially Adept Technologies, and in 2004 on Trust in the Ambient Society. The proposed workshop is a follow-up from the latter.

Pam Briggs holds the Chair in Applied Cognitive Psychology at Northumbria University. She is the Research Director for the Psychology and has recently been involved in the formation of a new Centre for Cognition and Communication. She has led numerous industry and government funded projects on the socio-cognitive aspects of computer-mediated communication, with the most recent focussed upon trust and privacy. She has given various invited papers on the theme of socially responsible developments in HCI and is writing a book on this topic. Her most recent project is a three year study which aims to model the process of seeking and evaluating health advice online – a project which integrates the research themes of trust, risk, self-disclosure and social identity. Pam has participated in several workshops, including recent CHI workshops on Socially Adept Technologies and Personalisation. She co-organised the 2004 CHI workshop on this topic.

Waleed Wagealla is a research fellow at the department of Computer and Information Sciences, the University of Strathclyde in Glasgow. He received BSc in computer science from University of Khartoum - Sudan, a PhD in computer science from Nottingham Trent University. His recent work is in mobile agents technology and trust management in pervasive computing systems. He has

published several papers in the fields of agent technology and trust management. He co-chaired/organized the first internal iTrust workshop, Glasgow, September 2002 and the special session on Trust Management in Collaborative Global Computing at the WETICE workshop on "Enterprise Security"', Linz, June 2003.

REFERENCES AND SELECTED READING

1. Bos, N., Olson, J., Gergle, D, Olson, G., Wright, Z., 2002,. Effects of four computer-mediated communications channels on trust development. CHI 2002: 135-140
2. Jones, S. and Marsh, S., 1997. Human-Computer-Human Interaction: Trust in CSCW. SIGCHI Bulletin, July 1997. 29(3):36-40.
3. Luhmann, N., 1979. Trust and Power. Chichester: Wiley.
4. Marsh, S., Briggs, P., Wagealla, W., 2004, Considering Trust in the Ambient Society, CHI 2004 Extended Abstracts.
5. Marsh, S., 1995, Exploring the Socially Adept Agent, in Proceedings DIMAS 95: Distributed Intelligent Multi Agent Systems, Poland.
6. Marsh, S., 1994, Formalising Trust as a Computational Concept., PhD Thesis, University of Stirling, Scotland.
7. Olson, G., & Olson, J. (2000). Distance matters. Human Computer Interaction, 15, 139-179.
8. Zheng, J., Veinott, E., Bos, N., , Olson, J. , Olson, G., 2002. Trust without touch: jumpstarting long-distance trust with initial social activities. CHI 2002: 141-146
9. Zolin, R., Hinds, P., Fruchter, R. & Levitt, R., 2002, Trust in Cross-Functional, Global Teams. CIFE Working Paper #67 April 2002. Stanford University.