

Introduction to this Special Issue* on
Revisiting and Reinventing Email

Steve Whittaker

Sheffield University

Victoria Bellotti

Palo Alto Research Center

Paul Moody

IBM Research

Steve Whittaker is a cognitive psychologist with interests in the application of cognitive science principles and methods to the design of systems for computer mediated communication, collaborative systems, information retrieval, and personal information management; he is Chair of Information Retrieval at University of Sheffield. **Victoria Bellotti** is a technology-focused ethnographer with interests in user-centered design and evaluation, task and personal information management, computer mediated communication, and collaborative work; she is a Principal Scientist and Manager of the Socio-Technical and Interaction Research Area at the Palo Alto Research Center. **Paul Moody** is a design researcher pursuing research, design, and development of collaborative business applications that improve worker satisfaction and business efficiency; he is a researcher in the Collaborative User Experience group at IBM Research in Cambridge, MA.

* Special Issue of *Human-Computer Interaction*, Volume 20 (2005), Numbers 1 and 2.

Revisiting and Reinventing Email

Email is the most successful computer application yet invented. It is used by millions of people to carry out their business each day. Over the last 10 years email has become ubiquitous in the home, too, and it is often the reason for purchasing a home computer (Kraut et al., 2000). It has changed the way that people work, and the ways that organizations operate; many types of collaborative work would be unthinkable without it (Sproull and Kiesler, 1991).

Various reasons have been put forward for email's success. Unlike face to face communication, its affordances free participants from the constraints of space and time – allowing senders and recipients to communicate at times and in places that are convenient to each (Clark and Brennan, 1991, Sproull and Kiesler, 1991, Kraut and Attewell, 1997). Another significant property is its *malleability*. Studies of email usage have repeatedly documented the striking number of different purposes to which it is put: email can support conversations, operate as a task manager, document delivery system, archive, and contact manager - to name but a few (Bellotti et al., 1993, Mackay, 1988, Whittaker et al., 2002a, Whittaker and Sidner, 1996). And at a technical level, it operates using a highly simple protocol.

Yet some of these same success factors contribute to the problems that are now endemic in email. Users complain about the sheer number of messages they receive, they complain about the number of irrelevant spam messages that they get sent, and they complain that email doesn't provide direct support for the tasks they use it for (Bellotti et al., 1993, Whittaker et al., 2002a, Venolia et al., 2001, Whittaker and Sidner, 1996). These are important problems, not simply as a technological challenge. Solving them could have significant effects on general productivity because so much knowledge work is now channeled through email.

Yet despite its success and these problems, one paradox is how little email has changed in the last few years. Early text based systems were replaced by graphical user interfaces, but aside from a few minor modifications (attachments, HTML integration, folders and address books), today's systems are remarkably similar to those introduced 15 years ago. The goal of this special issue is not to address this paradox, but rather to present work that begins to tackle email's problems. We start with a short summary of what research has been done over those last 15 years to set the stage for the contributions of this special issue, and conclude with some speculation about what we see as being important future trends and questions.

If we are to design better systems we need to know how people currently use email and what problems they experience in using it. And one research area has been *empirical studies of email usage*. For example, research has looked into people's problems with: *filing* (Bälter 2000, Whittaker and Sidner, 1996), *task management* (Bellotti et al., 1993,

Whittaker and Sidner, 1996), *triage* – i.e. deciding what to do with incoming messages (Bellotti et al., 2003, Venolia et al., 2001), *contact management* (Whittaker et al., 2002b), *semi-structured messaging* (Mackay et al., 1989), and *individual differences* (Gwidzka, 2000, Whittaker and Sidner, 1996). However one problem with this research is the shortage of data. In many of these areas, one or two pioneering studies have addressed a specific user problem, but we lack systematic information - information that would be highly informative in the generation of new designs. It is ironic to contrast this situation with the massive empirical literature on organizational uses of email (Sproull and Kiesler, 1991), or media richness (Daft and Lengel, 1984, see Walther, 1996, for a review). Here there is much data, but it is unclear how it can directly inform system design.

A second research area has been *novel system designs*. Here again there have been multiple approaches. Some designs have tackled the problem of *task management* either by providing direct support for tasks (Bellotti et al., 2003, Whittaker et al., 1997) or by thread visualization to help users detect relations between related Inbox items (Rohall et al., 2001, Venolia et al., 2001). A related line of research treats email as a form of *workflow* and proposes interfaces that overlay workflow support onto an email infrastructure (Borenstein and Thyberg, 1988, Borenstein, 1992). Other research addresses the *filing* problem using machine learning techniques to classify Inbox messages (Segal and Kephart, 1999). Machine learning researchers have also viewed email as an attentional problem and used similar learning techniques to classify incoming messages according to their *importance* (Horvitz, 2003). There have also been attempts to design email interfaces that reflect *different cognitive styles* (Gwidzka, 2000) or activity levels (Takkinen, and Shahmehri, 1999). More recently several researchers have proposed people-centric *social interfaces* to email (Farnham, 2002, Fisher and Dourish, 2004, Whittaker et al., in press). Another potentially important trend might be the move towards *search-based* clients such as Gmail (<https://gmail.google.com>) by Google. Again a major problem with this design research lies in generalizing from these results; in most areas, one or two systems have been built and tested, and it is still not obvious which approaches are promising and which are not.

While building effective systems requires careful empirical analysis and design, it is also clear that designs can be informed by theoretical guidelines or perspectives. Yet there has been conspicuously little research into *theory*. Part of the problem here is the *malleability* of email. If we view email as communication tool, then we might want to apply theories of asynchronous communication. But if we view it as an archive we might apply search or digital library frameworks. And if we regard email as a collaboration tool, then we might apply workflow analyses. Finally if we view the email problem as one of attention allocation then we might apply attentional accounts. However even more so than for empirical studies and new designs, there is a real absence of theorizing about what email is and what it does. In our view this is a significant gap that needs to be addressed.

This special issue bears on all three research areas, with a major focus on empirical studies and new designs.

Ducheneaut and Watts set the stage by presenting a detailed review of email research - identifying three major metaphors that have guided past research. They argue that email has variously been regarded as: a *filing cabinet* – leading research to focus on questions about how people manage and access email archives; a *production line* – where the central questions concern how people collaborate with others to accomplish work tasks; and a *communication genre* – addressing issues about how email is used in organizational settings to carry out group work. They conclude by attempting to draw these three threads together in a generative design framework focused on communicative, individual and socio-organizational factors. As we noted above, there is a distinct lack of theorizing about email and this paper is a notable exception in drawing together and organize a large body of disparate literature.

The next three papers all follow a classic HCI design cycle: they all begin with empirical user studies, using these to derive design guidelines from which they develop and evaluate new email systems.

Whittaker combines empirical work, design principles and systems development. It addresses the problem of collaborative task management in email. Whittaker summarizes empirical work showing that users experience severe difficulties in using the Inbox to manage collaborative tasks. In particular, users have problems reminding themselves about outstanding tasks and in collating prior task context for extended collaborative tasks. Whittaker then presents two systems, TeleNotes and ContactMap, that directly address collaborative task management. Their designs are based on strategies motivated by empirical research into *paper-* and *people-based* task management. Whittaker describes how the systems implement these different task management strategies and presents evaluation data for each system in use. He contrasts the success of these two approaches with earlier systems approaches and discusses outstanding design and theory problems arising from the research.

Bellotti, Ducheneaut, Howard, Smith, and Grinter address a similar problem (which they dub interdependent task management) but arrive at a different design. They begin with a highly detailed set of empirical studies that attempt to identify users' main problems with email. They look at users' strategies for processing incoming email, critiquing prior taxonomies of processing strategies and developing a novel metric for perceptions of overload. This empirical work leads them to identify six main design requirements for email. They then present a novel system, TaskMaster that embeds task management directly into the client, as well as providing various features for managing task deadlines. They evaluate the system in a field trial, with positive results.

Wattenberg, Rohall, and Gruen tackle a related problem but situate their research in an organizational context. They summarize an extensive research program carried out at IBM, Cambridge. Early work here focused on mailbox usage, identifying individual differences in email processing strategies as well as collaborative uses of the Inbox. But their main focus is on attempts to visualize message traffic and message threads. A key

discussion concerns the tension between research focusing on radical technology development, and the demands of developing techniques that can be incorporated into existing enterprise email systems. The authors describe various promising techniques that were vetoed by product groups because of the constraints placed by existing system requirements. Wattenberg et al. conclude with some requirements for successful technology transfer.

Schmandt and Marti address what promises to be an important future issue: heterogeneous multidevice access to email. In particular they focus on mobile ubiquitous settings. They point out that an increasing trend is towards accessing email on mobile devices in situations where the user is often busy with other activities. This places a strong requirement on the development of filtering rules to ensure that only important messages reach the recipient. Other demands here are to track the various potential user access devices to determine which is the critical device to deliver important information. The authors describe an architecture which supports the development of these rules and a long-term implementation of their system, Active Messenger, analyzing their experiences in accessing email in this way.

Finally **Kraut, Sunder, Telang, and Morris** explore a radically different approach to email overload. They point out that email is unusual compared with other communication media (such as postal mail or advertising) because it is free to send a message. They present a laboratory study in which they explore user behaviors when senders have to pay to send a message. They find that differential pricing affects user willingness to send messages (suggesting that this might be one way to curtail spam), but also that recipient's reading preferences are not affected by how much the sender paid to post a message.

These papers also highlight new developments that are likely to transform email over the next few years. Three areas where we expect to see rapid developments are in *mobile access*, *text processing* and *task management*.

Mobility. Mobile phones and PDAs are becoming important platforms for many users, but they raise significant HCI challenges. First new interfaces have to be designed to finesse a lack of display space. This may serve to introduce new designs that break away from current standards. The demands of accessing information on the move may also mean that mobile workers will also expect more assistance from systems in identifying important messages for immediate processing, opening the door to filtering and personalization techniques, as well as interfaces that directly support mobile access.

Text Processing. The last 10 years since the appearance of the web have seen a transformation in techniques for text processing, e.g. novel information retrieval algorithms, information extraction techniques, and text classification methods. So far these have been applied to public corpora such as the web, or digital libraries, but recently researchers and companies have begun to apply them to personal information such as email (Cohen, 1996, Horwitz, 2003, Sahami et al., 1998). While these automatic techniques have great potential for addressing pressing problems with email, e.g. spam detection, it is also clear that there are significant issues associated with their integration into working email interfaces. Early data indicates that users are often mistrustful of such

techniques - having very high expectations and being somewhat intolerant of errors (Pazzani, 2000). It will significantly test the ingenuity and design skills of HCI researchers if we are to successfully integrate these new techniques into email UIs.

Task Management. A final area where progress is needed is in redesigning email to help people better manage their tasks. While novel text processing algorithms might help users better detect and cluster related messages in their Inboxes, there is also a more significant challenge; to design systems that support users in detecting and processing messages associated with important tasks. Currently users are at the mercy of an undifferentiated Inbox in which all messages have equivalent status. We need to devise systems that help users identify important messages quickly so that they can proactively focus on their important tasks, rather than sifting through multiple messages attempting to determine how each message might relate to their outstanding tasks.

NOTES

Acknowledgments. Thanks to Marilyn Walker for comments on early drafts.

Support. This research was partially funded under European Union Framework 6 grant *FP6-2002-IST-1*.

Editors' Present Addresses. Steve Whittaker, Department of Information Studies, University of Sheffield, 211 Portobello, Sheffield, S1 4DP, UK. Email: s.whittaker@shef.ac.uk. Victoria Bellotti, Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304. Email: bellotti@parc.com. Paul Moody, Collaborative User Experience Group, IBM Research, 1 Rogers Street, Cambridge, MA 02142. Email: paul_moody@us.ibm.com.

REFERENCES

- Bälter, O. (2000): Keystroke-level analysis of email message organization. *Proceedings of the CHI 2000 Conference on Human Factors in Computer Systems*, 105-112. New York: ACM Press.
- Bellotti, V., Duchenaut, N., Howard, M., and Smith, I. (2003). Taking email to task: the design and evaluation of a task management centered email tool. *Proceedings of the CHI 2003 Conference on Human Factors in Computer Systems*, 345-352. New York: ACM Press.
- Borenstein, N. (1992). Atomicmail. Computational Mail as a network infrastructure for computer supported co-operative work *Proceedings of the CSCW 1996 Conference on Computer Supported Co-operative Work*, 67-74. New York: ACM Press.
- Borenstein, N., and Thyberg C. (1988). Co-operative Work in the Andrew message system. *Proceedings of the CSCW 1996 Conference on Computer Supported Co-operative Work*, 306-323. New York: ACM Press.
- Clark H. & Brennan, S. (1991). Grounding in communication. In L.B. Resnick, J. Levine & S. Teasley, Eds. *Perspectives on socially shared cognition*. Washington D.C., APA Press.
- Cohen, W. (1996): Learning Rules that Classify E-Mail in *AAAI Spring Symposium on Machine Learning and Information Retrieval*, 18-25, Menlo Park, AAAI Press.
- Daft, R.L. & Lengel, R.H. (1984). Information richness: a new approach to managerial behavior and organizational design. In: Cummings, L.L. & Staw, B.M. (Eds.), *Research in Organizational Behavior* 6, (pp191-233). Homewood, IL: JAI Press.
- Farnham, S., Turski, A., Portnoy, W., & Davis, J. (2002). Connections: Exploring Who Knows Whom through Social Networks. *Paper presented at Human Computer Interaction Conference*, Winter Park, Colorado.
- Fisher, D. and Dourish, P. (2004). Social and Temporal Structures in Everyday Collaboration. *Proceedings of the CHI 2004 Conference on Human Factors in Computer Systems*. New York: ACM Press.
- Gartner. (2001). Trends in email usage.
- Gwizdka, J. (2002). Reinventing the Inbox. Supporting Task Management of Pending Tasks in Email. *Proceedings of the CHI 2002 Conference on Human Factors in Computer Systems*, 550-551. New York: ACM Press.

- Horvitz, E., Kadie, C. Paek, T., and Hovel, D. (2003). Models of attention in computing and communication, *Communications of the ACM*, 46(3), 52-59.
- Kraut, R. E. & Attewell, P. (1997). Media use in a global corporation: Electronic mail and organizational knowledge. In S. Kiesler (Ed.) *Culture of the Internet*. (pp. 323-342). Mahwah, NJ: Erlbaum.
- Kraut, R., Mukhopadhyay, T., Szczypula, J., Kiesler, S., & Scherlis, B. (2000). Information and communication: Alternative uses of the Internet in households. *Information Systems Research*, 10, 287-303.
- Levitt, M. (2000). Email usage forecast and analysis, 2000-2005, IDC Report W23011.
- Mackay, W. (1988): More Than Just a Communication System: Diversity in the Use of Electronic Mail, in *Proceedings of the CSCW 1998 Conference on Computer Supported Co-operative Work*, 344-353. New York: ACM Press.
- Mackay, W. E., Malone, T. W., Crowston, K. Rao, R., Rosenblitt, D., & Card, S. K. (1989) How do experienced information lens users use rules?, *Proceedings of the CHI 1989 Conference on Computer Supported Co-operative Work*. 211-216. New York: ACM Press.
- Pazzani, M. (2000). Representations of email filtering profiles: a user study. *Proceedings of IUI 2000 Conference on Intelligent User Interfaces*, 202-206. New York: ACM Press.
- Pitney Bowes (2000). Increased use of electronic communications tools among North American and European workers, press release, August, 2000.
- Rohall, S., Gruen, D., Moody, P., Kellerman, S. (2001). Email visualizations to aid communications. *Proceedings of IEEE Symposium on Information Visualization*, 12-15. New York: IEEE Press.
- Sahami, M., Dumais, S., Heckerman, D., and Horvitz, E. (1998). A Bayesian Approach to Filtering Junk E-Mail. In AAI Workshop on Learning for Text Categorization, 55-62, Madison Wisconsin. AAI Technical Report WS-98-05.
- Segal, R., and Kephart, J. (1999). MailCat: an intelligent assistant for organizing email. *Proceedings of Conference on Autonomous Agents*, 276-282, ACM Press. New York.
- Sproull, L., and Kiesler, S (1991). *Connections*. Cambridge, MIT Press.
- Takkinen, J., and Shahmehri, N. (1999). Task-oriented restructuring of an application domain: a multi-agent architecture for doing things in internet email. *Proceedings of HICCS-32*. New York: IEEE Press.

- Venolia, G., Gupta, A., Cadiz, J.J., and Dabbish, L. (2001). Supporting Email Workflow. (Microsoft Technical Report. MSR-TR-2001-88).
- Walther, J. B. (1996). Computer-mediated communication: Impersonal, interpersonal, and hyperpersonal interaction. *Communication Research*, 23, 1-43.
- Whittaker, S. and Sidner, C. (1996). Email Overload: Exploring Personal Information Management of Email. *Proceedings of CHI 96 Conference on Human Factors in Computing Systems*, 276-283. New York: ACM Press.
- Whittaker, S. Jones, Q., and Terveen, L. (2002a). Persistence and Conversation Stream Management: Conversation and Contact Management. *Proceedings of HICCS'02*, 335-344. New York: IEEE Press.
- Whittaker, S. Jones, Q., and Terveen, L. (2002b). Contact Management: Identifying Contacts to Support Long term Communication, *Proceedings of CSCW 2002 Conference on Computer Supported Cooperative Work*, 216-225. New York: ACM Press.
- Whittaker, S., Jones, Q., Nardi, B., Creech, M., Terveen, L., Isaacs, E., Hainsworth, J. (in press). ContactMap: Organizing Communication in a Social Desktop. To appear in *ACM Transactions on Computer Human Interaction*.
- Whittaker, S., Swanson, G., Kucan, J., & Sidner, C., (1997). Telenotes: managing lightweight interactions in the desktop. *Transactions on Computer Human Interaction*, 4, 137-168.

ARTICLES IN THIS SPECIAL ISSUE

- Bellotti, V., Duchenaut, N., Howard, M., Smith, I., and Grinter, R. E. (2005). Quality versus quantity: Email-centric task-management and its relationship with overload. *Human-Computer Interaction*, 20, nnn-*nnn*.
- Duchenaut, N., and Watts, L. A. (2005). In search of coherence: A review of email research. *Human-Computer Interaction*, 20, nnn-*nnn*.
- Kraut, R. E., Sunder, S., Telang, R., and Morris, J. (2005). Pricing electronic mail to solve the problem of spam. *Human-Computer Interaction*, 20, nnn-*nnn*.
- Schmandt, C., and Marti, S. (2005). Active messenger: Email filtering and delivery in a heterogeneous network. *Human-Computer Interaction*, 20, nnn-*nnn*.
- Wattenberg, M., Rohall, S. L., Gruen, D., and Kerr, B. (2005). Email research: Targeting the enterprise. *Human-Computer Interaction*, 20, nnn-*nnn*.
- Whittaker, S. (2005). Supporting collaborative task management in email. *Human-Computer Interaction*, 20, nnn-*nnn*.