Connecting versus Calming: Interruptibility, Presence, and Availability

Position Paper for the CHI 2004 Workshop on Forecasting Presence and Availability

James Fogarty
Human Computer Interaction Institute
Carnegie Mellon University
jfogarty@cs.cmu.edu

Many researchers in the HCI and CSCW communities have examined the problem of making it easier for people to connect. Work on media spaces, for example, examined the role of always-on audio and/or video connections in supporting lightweight encounters [1, 4, 6], together with the distractions and privacy concerns associated with such connections [15, 17].

Recently, attention has focused on providing better support for commonly used communication tools, such as email, instant messaging, and the telephone. These tools each have their own strengths, but as a whole generally do not support awareness of a colleague's actions. In contrast, face-to-face encounters and media spaces both allow a person to gather information about a person's context prior to attempting to initiate communication. Work has sought to enhance these communication tools by examining patterns of presence [2, 3] and patterns of attendance at meetings [12, 16]. Included in interfaces, this sort of information allows people to better estimate when they might be able to communicate with a colleague. Other work has examined patterns of presence to make decisions about whether to forward urgent email messages to a person's mobile phone [10].

Much of this prior work is focused on enabling easier connections, based on the viewpoint that a person's absence is an obstacle to communication that needs to be overcome. While this is certainly the correct interpretation of the problem for some communication, there also seems to be a need for a calming influence on our communication. If a tool focuses only on making it easier to reach another person, that tool is neglecting another important aspect of face-to-face communication, our ability to recognize that a colleague is currently too busy to be interrupted and that we should instead defer our desire to communicate. In an indication that current tools do not currently allow such a determination, Hudson *et al* found that research managers reported physically moving away from their computers, or even away from their offices, in order to have uninterrupted working time [13].

This position paper reviews our work on sensor-based statistical models of human interruptibility and how such models might be used to better support computer-mediated communication. While participating in this workshop, I intend to focus on the idea that our attention should be focused not only on whether it is *possible* to communicate with a person, but also on whether it is currently *appropriate*.

In a series of studies [7, 8, 14], we have explored the feasibility and the robustness of sensor-based statistical models of human interruptibility. Examining office workers in their natural working environments, we collected randomly timed interruptibility self-reports. We then considered a variety of sensors, in combination with widely-used machine learning techniques, to determine which sensors provided the most reliable estimates of interruptibility.

The initial results in this line of work were based on audio and video recordings. Human subjects using the recordings of an unknown person working in their office were able to distinguish between self-reported "Highly Non-Interruptible" situations and other situations with an accuracy of 76.9%. Statistical models based on simulated sensors in the audio and video recordings were able to make this same distinction with an accuracy as high as 82.4%, significantly better than the human performance. Interestingly, much of the accuracy of the statistical models came from only a few sensors. By itself, a simulated sensor to determine whether anybody in an office was talking had an accuracy of 75.9%, not significantly different from the human performance. Combining this simulated talking sensor with simulated sensors for keyboard or mouse activity, for using the phone, and for the time of day, we built models with an accuracy of 79.2%, not significantly different from the human performance.

We have since validated these results using real sensors over a larger and more diverse group of office workers. Our approach was able to reliably create models with accuracies exceeding 80%, as good as or better than the human performance. We have also had positive results in examining interesting subsets of sensors, including combinations that could be deployed entirely in software with no infrastructure costs and combinations that could be sensitive to certain privacy concerns by using neither microphones nor cameras. These successes give us reason to believe that reliable statistical models can be created for a variety of office workers in a broad set of circumstances.

In work related to these studies [9], we have examined one approach to using such models in computer-mediated communication. Using a socially translucent computing approach [5], we built a context-aware communication client that uses a model of interruptibility to show colleagues how busy a person currently is. This system shares this information with colleagues, but leaves to people the decision on when and how to communicate (as opposed to an automatic filtering approach, like that taken by the Notification Platform [11]). We then deployed this system with four groups of colleagues to see how the provided interruptibility and context information affected their decisions to communicate.

Somewhat surprisingly, users of our system primarily used context as an indication of presence. While we have shown that talking in an office is a very good indicator that a person is currently busy and our client shows colleagues as unavailable if talking is detected, the trend among participants in our study was actually to be more likely to initiate an instant messaging session when a colleague was shown as talking ($\chi2(1, 1411) = 1.28, p \approx .26$). Though the difference is not significant, we had expected the difference to be significant in the other direction. Participants were, however, significantly less likely to initiate an instant messaging session with somebody whose computer had been inactive for several minutes. These findings would seem to be consistent with people making decisions based on their own availability and their own desire to communicate, rather than considering a colleague's interruptibility.

These findings raise questions about how to include interruptibility and context information in computer-mediated communication tools. One possibility is that programs might share context information when a person is interruptible, but become stingy with context information when a person is non-interruptible. This is similar to the strategy that Hudson *et al* found among research managers, who used their presence to regulate their interruptibility. An interruptible person might therefore be shown as currently in their office, but it would not be clear whether a non-interruptible person had left their office or was engaged in a conversation in their office. Such a design might encourage communication when it was appropriate for both parties.

References

1. Adler, A. and Henderson, A. (1994) A Room of Our Own: Experiences from a Direct Office Share. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1994)*, 138-144.
2. Begole, J.B., Tang, J.C. and Hill, R. (2003) Rhythm Modeling, Visualizations, and Applications. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2003)*.
3. Begole, J.B., Tang, J.C., Smith, R.B. and Yankelovich, N. (2002) Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2002)*, 334-343.
4. Bly, S.A., Harrison, S.R. and Irwin, S. (1993) Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment. *Communications of the ACM*, *36* (1). 28-46.
5. Erickson, T. and Kellogg, W.A. (2000) Social Translucence: An Approach to Designing Systems that Support Social Processes. *ACM Transactions on Computer-Human Interaction (TOCHI)*, *7* (1). 59-83.
6. Fish, R.S., Kraut, R.E., Root, R.W. and Rice, R.E. (1993) Video as a Technology for Informal Communication. *Communications of the ACM*, *36* (1). 48-61.
7. Fogarty, J., Hudson, S., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. and Yang, J. (2003) Predicting Human Interruptibility with Sensors. *In Review, ACM Transactions on Computer-Human Interaction (TOCHI)*.
8. Fogarty, J., Hudson, S. and Lai, J. (2004) Examining the Robustness of Sensor-Based Statistical Models of Human Interruptibility. *To Appear, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2004)*.
9. Fogarty, J., Lai, J. and Christensen, J. (2004) Presence versus Availability: The Design and Evaluation of a Context-Aware Communication Client. *To Appear, International Journal of Human-Computer Studies (IJHCS)*.
10. Horvitz, E., Jacobs, A. and Hovel, D. (1999) Attention-sensitive alerting. *Proceeding of the Conference on Uncertainty and Artificial Intelligence (UAI 1999)*, 305-313.
11. Horvitz, E., Kadie, C., Paek, T. and Hovel, D. (2003) Models of Attention in Computing and Communication: From Principles to Applications. *Communications of the ACM*, *46* (3). 52-59.
12. Horvitz, E., Koch, P., Kadie, C.M. and Jacobs, A. (2002) Coordinate: Probabilistic Forecasting of Presence and Availability. *Proceedings of the National Conference on Uncertainty and Artificial Intelligence (UAI 2002)*.
13. Hudson, J.M., Christensen, J., Kellogg, W.A. and Erickson, T. (2002) "I'd be overwhelmed, but it's just one more thing to do": Availability and Interruption in Research Management. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2002)*, 97-104.
14. Hudson, S., Fogarty, J., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. and Yang, J. (2003) Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2003)*, 257-264.
15. Hudson, S.E. and Smith, I. (1996) Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 1996)*, 248-257.
16. Mynatt, E. and Tullio, J. (2001) Inferring Calendar Event Attendance. *Proceedings of the International Conference on Intelligent User Interfaces (IUI 2001)*, 121-128.
17. Smith, I. and Hudson, S.E. (1995) Low Disturbance Audio for Awareness and Privacy in Media Space Applications. *Proceedings of the ACM International Conference on Multimedia*, 91-97.