

Buzz: Measuring and Visualizing Conference Crowds

Christopher Wren, Yuri Ivanov, Darren Leigh

TR2007-035 March 2008

Abstract

This report is a proposal for an exhibition that will explore the idea of using technology to understand the movement of people. Not just on a small stage, but in an expansive environment. Not the fine details of movement of individuals, but the gross patterns of a population. Not the identifying biometrics, but the group behavioral patterns that evolve from the structure of the environment and the points of interest embedded in that structure. In this instance: a marketplace, and in particular, the marketplace of ideas called Emerging Technologies at Siggraph 2007.

ACM Siggraph 2007 (Emerging Technologies), August 2007

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

MITSUBISHI ELECTRIC RESEARCH LABORATORIES
<http://www.merl.com>

Buzz: Measuring and Visualizing Conference Crowds

Christopher R. Wren Yuri A. Ivanov Darren Leigh
Jonathan Westhues

TR2007-035 August 2007

Abstract

This report is a proposal for an exhibition that will explore the idea of using technology to understand the movement of people. Not just on a small stage, but in an expansive environment. Not the fine details of movement of individuals, but the gross patterns of a population. Not the identifying biometrics, but the group behavioral patterns that evolve from the structure of the environment and the points of interest embedded in that structure. In this instance: a marketplace, and in particular, the marketplace of ideas called Emerging Technologies at SIGGRAPH 2007.

Appears in SIGGRAPH Emerging Technologies 2007

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Copyright © Mitsubishi Electric Research Laboratories, Inc., 2007
201 Broadway, Cambridge, Massachusetts 02139

August 2007: Invited participation in the SIGGRAPH 2007 Emerging Technologies demo section.
<http://doi.acm.org/10.1145/1278280.1278307>




Figure 1: The wireless, low-power sensor used in the system.

1 Introduction

This exhibition explores the idea of using technology to understand the movement of people. Not just on a small stage, but in an expansive environment. Not the fine details of movement of individuals, but the gross patterns of a population. Not the identifying biometrics, but patterns of group behavior that evolve from the structure of the environment and the points of interest embedded in that structure. In this instance: a marketplace, and in particular, the marketplace of ideas called SIGGRAPH 2007 Emerging Technologies (E_{TECH}).

In this exhibit we combine sensor hardware, analytic tools, and visualization technology into a system to answer some of the most basic questions about a trade show venue like E_{TECH}: where are the people right now? which exhibits always have a crowd? when is the venue most crowded? and how do other events such as keynotes affect venue traffic? We cannot say that we truly understand the value of the venue, if we cannot even answer these basic questions about how the participants utilize the space.

The first step in understanding a thing is observing it. At the Mitsubishi Electric Research Laboratories (MERL) we have developed a sensor network research platform specifically tuned to observing the activities of people in a built environment. A sensor node, without its enclosure, is pictured in Figure 1. The nodes are in some ways descended from the work at M.I.T/ [10] but they include a number of improvements such as power efficient design for multi-year battery life, a standards-based MAC for communications reliability, and a design more suitable to large-scale manufacturing to support this kind of large-scale project.

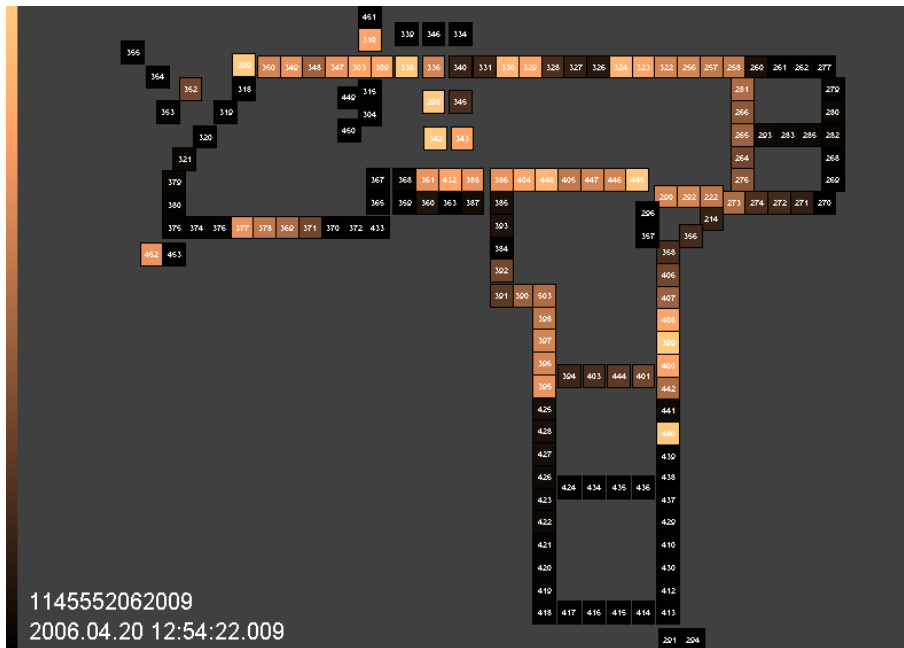


Figure 2: An example of high-population movement during an evacuation.

This exhibit builds on the pioneering work in sensor networks in general such as Estrin [4] and the Berkeley Motes team [5]. We choose a dense network of simple sensors over a sparse collection of cameras [11] for the several reasons: cost, robustness, and privacy. The network is much more cost effective. Of course the sensors themselves are very inexpensive, but also the simplicity of the data stream generated by the sensor leads to lower computational overhead and lower communication overhead. Those both imply lower power consumption. The nodes pictured in Figure 1 draw an average of $200\mu W$. This means that the sensors can be wireless and run off battery or parasitic power sources. This means that installation can be extremely simple and fast, drastically lowering the overall cost of the system. Several groups have explored the benefits of networks of motion sensors for monitoring human behavior, primarily in residential settings [2, 13, 10, 1].

A network of simple sensors is more robust in the sense that single node failures do not have as significant an impact on the data stream as it would in the case of high-value nodes. Since the algorithms and hardware of each node are significantly simpler than what would be required in a smart camera, with integrated computer vision, possibly compression, and high-bandwidth communication, they can be engineered to a higher level of robustness to environmental, situational, and algorithmic failures.

The sensor modality itself is inherently privacy-friendly in the sense that it records far less information than a camera: only about one bit per square

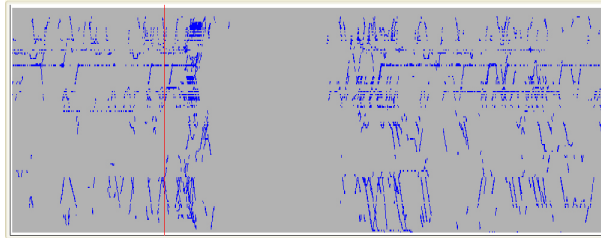


Figure 3: The timeline view of the evacuation, showing both gross and fine behavior features.

meter per second. Not enough to identify individuals, certainly, or even enough to detect any but the most overt behaviors or moving about the environment. With a camera-based system the system must be carefully constructed to scrub privacy-invading information from the collected data, and the observed can only trust that the system designers and operators have done that job. With motion sensors the information is never sensed, and so there is no need for the observed to trust the designer or operator at all[12]

All that remains is to prove that the networks can be useful as well as affordable, robust, and sensitive to human needs.

2 Sensation on the Network

Figure 2 shows the response of a network of motion detectors to a flood of evacuation activity during a false fire alarm at our research facility in Cambridge. The map shows the geographic distribution of 155 sensors covering the hallways, lobbies, and public meeting places on one floor of the office building where we work. The brighter the sensor the more recently it was activated, with black sensors not having been activated within approximately the last 10 seconds.

Despite having about 100 inhabitants (significantly more in the summer), this system typically only observe a handful of people moving around the space at any given time. Certain events, such as the fire evacuation in Figure 2 are the exception: some trigger causing the synchronous movement of a great many people, driving the network toward saturation. However, it is still possible to see significant structure in the data. For example, there are many people traveling from the seminar room in the upper left of the map to an emergency exit near the center of the building. Even in this visualization the signature of this massive flow of people is noticeably different from the movements of single individuals and small groups evident in the rest of the space at that time. It's amusing to note that the path chosen by the throng was in fact erroneous: there is a different exit that is much closer to the seminar room that should have been used instead.

A different way to view the data is pictured in Figure 3: a timeline or piano-roll[6] mode that reveals temporal patterns not apparent from the spatial view.




Figure 4: Another view showing a whole week, with gross daily and weekly behavioral patterns in evidence.

The time axis is along the horizontal, and the sensors are stacked up on the vertical axis. It is easy to see several distinct states of the building: normal activity, heightened activity during evacuation, and the empty post-evacuation state. At a larger time-scale in Figure 4, we see an entire week of data. Gross patterns become apparent at this scale: more activity during the day compared to at night, and more activity on the weekdays compared to weekends. It is even apparent that activity begins and ends at different times of day in different parts of the space.

3 Perception on the Network

Visualizing raw sensor data can be very powerful, as it can provide immediate insight into complex patterns of behavior. Often these patterns are invisible to the participants. It is difficult to get the Gestalt of these patterns while one is down in the them, moving around with the rest of the ants, to use an apt, if unkind metaphor. We have developed a toolbox of analytic techniques that work well on these kinds of data streams. The analytic techniques allow us to extract the more subtle patterns and expose them in concisely focused visualizations.

Figure 5 shows the power of such analysis. The system uses statistical models of activity to extract form from the mass of raw activations. The top plot shows the spatial distribution of motion energy in the space: the parts of the space that see the most motion activations are colored in dark. The bottom plot tells a very different story by extracting only the activities that appear to be related to meetings. This washes away all the activations that have to do with walking, or loitering, or any of the other varied activities in the space. Given this analysis, it is possible to visualize a use pattern that is invisible from the raw activations alone.

4 Interacting with the Network

We expect these basic visualization and analytic tools to carry directly into the ETECH venue. Figure 6 shows a touch-table interface powered by the Diamond-Touch technology [3]. This interface is optimized for surveillance tasks [7] in office or warehouse environments where there are relatively few people moving around in an uncoordinated fashion.

The phrase “a few uncoordinated people” does not describe the ETECH venue particularly well. The methods presented above will need to be adjusted to deal well with the throngs of individuals loitering and browsing their way through ETECH. We have gathered data at a university research open house that we will use to further tune our algorithms for new behaviors and denser populations.

ETECH is also more like a marketplace of ideas [9]. The things people want to learn from the system are also very different. The participants will be asking questions of a more social nature [8]. A rushed participant might want a list of the most popular exhibits. Someone with more time might want to know where the crowds are right now, so that they can avoid those places and see the exhibits in more depth. Someone who wants to come back later might want to know, based on yesterday’s data, what time the exhibits are least congested. Conversely a photographer might want to know the best time to find that perfect crowd for an exciting photo.

The exhibit will present a transformed interface that is more visually attractive and compellingly interactive. It will allow visitors to navigating the ebbs and flows of the ETECH venue over the course of the week. We also plan to allow people to query the abstracts of various exhibits so that they can use this information in context, with the throng and popularity visualizations to plan their visits more effectively.

Acknowledgments

This work is funded by the Mitsubishi Electric Corporation and is made possible by the foresight and wisdom of research directors at MERL, both past and present. We thank Ishwinder Kaur of the Massachusetts Institute of Technology for providing the throng data from the Media Laboratory Digital Life meeting. We are grateful to Dr. Joseph Marks, Prof. John Sibert, and Dr. Kathy Ryall for inviting us to participate in the Emerging Technologies venue this year.

References

- [1] Gregory Abowd, Aaron Bobick, Irfan Essa, Elizabeth Mynatt, and Wendy Rogers. The aware home: Developing technologies for successful aging. In *Proceedings of AAAI Workshop on Automation as a Care Giver*, 2002.
- [2] Ryan Aipperspach, Elliot Cohen, and John Canny. Modeling human behavior from simple sensors in the home. In *Proceedings Of The IEEE Conference On Pervasive Computing*, 2006.

- [3] Paul H. Dietz and Darren L. Leigh. Diamondtouch: A multi-user touch technology. In *Symposium on User Interface Software and Technology (UIST)*, pages 219–226. ACM, November 2001. also MERL TR2003-125.
- [4] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. Next century challenges: scalable coordination in sensor networks. In *MobiCom '99: Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 263–270, New York, NY, USA, 1999. ACM, ACM Press.
- [5] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David E. Culler, and Kristofer S. J. Pister. System architecture directions for networked sensors. In *Architectural Support for Programming Languages and Operating Systems*, pages 93–104, 2000.
- [6] H. B. Horton. Machine for registering music, 1855. US Patent Office 13,946.
- [7] Yuri Ivanov and Christopher Wren. Toward spatial queries for spatial surveillance tasks. In *Pervasive: Workshop on Pervasive Technology Applied to Real-World Experiences with RFID and Sensor Networks*, volume EI123, May 2006.
- [8] Ishwinder Kaur. Openspace: Enhancing social awareness at the workplace. Master's thesis, Massachusetts Institute of Technology, 2007. in submission.
- [9] Steve A. Miller. *How to Get the Most Out of Trade Shows*. McGraw-Hill Professional, 2000.
- [10] E. Munguia Tapia, S. S. Intille, L. Lopez, and K. Larson. The design of a portable kit of wireless sensors for naturalistic data collection. In *Proceedings of PERVASIVE 2006*, Dublin, Ireland, 2006. Springer-Verlag.
- [11] Ali Rahimi, Brian Dunagan, and Trevor Darrell. Simultaneous calibration and tracking with a network of non-overlapping sensors. In *Computer Vision and Pattern Recognition*, pages 187–194. IEEE Computer Society, June 2004.
- [12] Carson J. Reynolds and Christopher R. Wren. Worse is better for ambient sensing. In *Pervasive: Workshop on Privacy, Trust and Identity Issues for Ambient Intelligence*, May 2006.
- [13] Daniel H. Wilson and Chris Atkeson. Simultaneous tracking & activity recognition (star) using many anonymous, binary sensors. In *The Third International Conference on Pervasive Computing*, pages 62–79, 2005.




Figure 5: **Top:** The spatial distribution of raw motion energy compared to the **Bottom:** spatial distribution of *meeting* activity detections




Figure 6: A data-exploration tool running on the DiamondTouch.