Geocentric Crowdsourcing and Smarter Cities: Enabling Urban Intelligence in Cities and Regions

A position paper for the 1st International workshop on ubiquitous crowdsourcing

Thomas Erickson
IBM T. J. Watson Research Center
snowfall@us.ibm.com

ABSTRACT

This paper explores the application of crowdsourcing to cities and regions – what I call "geocentric crowdsourcing." My basic interest is in how crowdsourcing can be used to enable cities and regions to more effectively address issues ranging from infrastructure to governance. I have two entwined aims in this position paper. One is to discuss the notion of smarter cities, and the ways in which crowdsourcing might contribute. The other is to suggest some distinctions for thinking about crowdsourcing: a four quadrant model of crowdsourcing, and computational vs. deliberative crowdsourcing. I intend these distinctions as grist for a workshop activity aimed at developing a principled way of thinking about crowdsourcing systems.

Author Keywords Crowdsourcing, geocompting, social computing, urban informatics, urban intelligence

INTRODUCTION: SMARTER CITIES

Sometime in 2009, according to demographers, the percentage of the world's population that resided in cities crossed the fifty percent mark. This, coupled with an array of increasingly severe environmental pressures, suggests that cities are an appropriate focus for R&D aimed at making cities more livable and sustainable places.

And, indeed, cities have been receiving increasing attention from purveyors of technology. Accenture has formed an "Intelligent City Network," to promote smart grid deployments. Cisco speaks of "Smart+Connected Communities." Hewlett-Packard talks of a "Central Nervous System for Earth." IBM includes "smarter cities" as a focal area in its "Smarter Planet" initiative. And so on.

These initiatives are all strongly focused on technology. The dominant vision of "smart technology" involves deploying networked sensors, gathering and aggregating their data, running analytics on it, and piping the results into control systems and dashboards overseen by managers (Fig 1). However, in these visions, people are typically relegated to the periphery. They are often treated as passive subjects, as the recipients of the benefits of "smart technology."

Over the last decade, an alternative, more social view of "smartness" has emerged. In this view, systems gain intelligence from both digital and human elements. People are seen as being actively engaged in systems as participants, and as a consequence can contribute their considerable knowledge and expertise to systems (Fig 2).

Just as sensors can gather information, so can people. Just as digital systems can analyze and model data, so can people. Just as control systems can act in the world, so can people. Of course the contributions people make differ from those of digital systems. Humans can contribute deep, qualitative knowledge; they can analyze fuzzy or incomplete data; and they can act in ways that digital systems often can not.

Besides offering the possibility of increasing the range of problems and quality of solutions, tapping human intelligence to augment the intelligence of cities offers the prospect that "smart systems" will be more likely to be accepted and viewed as legitimate. Rather than people as passive subjects of increasingly 'smart' technical systems, the vision is that smarter cities can offer a variety of ways for humans to act as first class participants, contributing their abilities to sense, analyze and act.

CROWDSOURCING FOR CITIES AND REGIONS

Since its coinage in 2006, the term "crowdsourcing" has become increasingly popular, and as a consequence its definition has become increasingly fuzzy. For the purposes of this position paper, I define "crowdsourcing" as "*the use of the perceptual and cognitive abilities of a large group of individuals to solve a problem*."

As noted, I am particularly interested in the application of crowdsourcing to cities and the regions that contain them. There are a number of reasons that crowdsourcing makes particular sense for urban contexts. First, that's where the people are – by definition, cities are dense concentrations of people. Second, inhabitants of a place develop a deep knowledge of it because they live, work and socialize there. Third, inhabitants have a practical interest in participating in systems that impact their daily life: someone who may never vote in an election, may still complain vociferously about potholes in a local street. Finally, inhabitants of a city or region often identify with it, or at least with networks of family, friends and communities associated with it; this pre-existing social structure is a valuable asset that other systems (e.g., Wikipedia) must develop from scratch.

Three Examples

Traditional crowdsourcing involves large scale systems that draw on the general knowledge of independent members of an anonymous crowd. What interests me is the prospect that crowdsourcing systems for cities and regions may take on a rather different character because of their focus on local knowledge and motivation of people situated in a particular place. Let's look at some examples.

Investigate Your MP

"Investigate Your MP" [9] is a site run by *The Guardian*, a London newspaper, in the wake of a scandal about excessive expense submissions. The site invited citizens to analyze expense reports submitted by MPs, and to flag those that deserved closer scrutiny by the site's administrators. (Currently 27,009 individuals have reviewed 221,351 of 458,832 pages, with a number of 'interesting' findings reported.) While this system does not tap into l*ocal knowledge* – it only requires commonsense knowledge about what are valid expenses – it does tap into *local motivation*, leveraging public indignation about the purported misbehavior of local representatives.

FixMyStreet

FixMyStreet [6] is an application that allows individuals to report potholes, streetlight outages, and other street-related problems on a publicly visible map (Figure 3). The problems are then brought to the attention of the appropriate council [governing body] responsible for fixing them. As individual's reports appear on the shared map, it creates a powerful aggregate representation of the state of the streets, and areas with lots of problems. FixMyStreet has been criticized, however, for poor integration with the government bodies responsible for addressing problems. Officials report frustration at being unable to note that a problem has been fixed, or is scheduled to be fixed, or can't be fixed due to a shortage of funds. [12]. One issue that this type of application raises (there are many similar ones such as FillThatHole and SeeClickFix) is how to manage the interaction between responsible officials and the crowd.

Figure 3. The FixMyStreet map.

Cyclopath

Cyclopath is a computational geo-wiki, a user-editable map that can be used to compute routes between points, said computations making use of the edited map data. Cyclopath is intended for bicyclists, and enables them to find bicycle-friendly routes around the city; it also relies on the cycling community to add data – road surface conditions, off-road paths, location of coffee shops – that is useful in determining a good bicycle route, but not found on conventional maps. Cyclopath has approximately 1500 registered users, who over the first 9 months of its use made edits to the map that resulted in routes that were approximately 8% (1K) shorter.

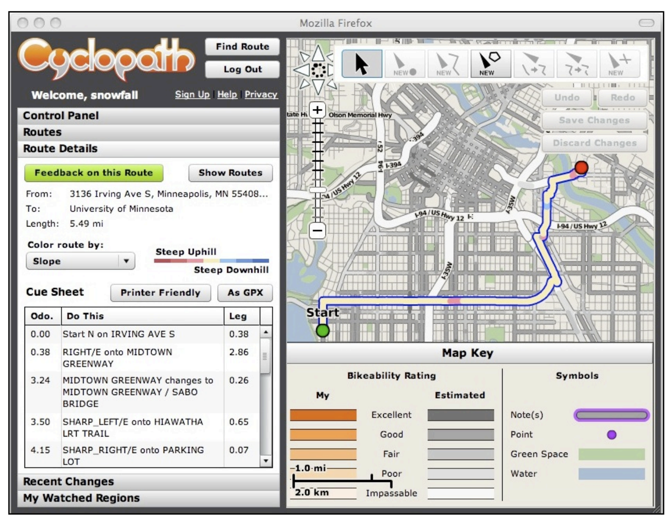

Figure 4. The Cyclopath geowiki.

An interesting aspect of Cyclopath can be seen by contrasting it with Google Map's recent introduction of bicycle routing. Google Maps offers bike routing as a generic service; in contrast, Cyclopath is local – it relies on a place-based community to contribute local knowledge, and its existence and use is a point of pride for the local bicycling community. To me, this gives Cyclopath, a very different feel – and a very different set of social dynamics – from systems like Google Map, without communities, or even Wikipedia, which has a community but where place is largely irrelevant.

TOWARDS A LANGUAGE FOR CROWDSOURCING?

Stepping back from the domain of cities and regions, I'd like to raise some more general issues. One thing that I would find useful – and that I hope might interest others in this workshop – is to explore dimensions useful for characterizing crowdsourcing systems. As grist for such a discussion, I offer two ways of thinking about crowdsourcing systems.

Four types of crowdsourcing

One way I think about crowdsourcing is in terms of the venerable four quadrant model from the early days of CSCW [10, 1]. This model divided cooperative work into quadrants based on its distribution over time and space: same time - same place; same time - different places; different times - same place; and different times - different places. Viewing crowdsourcing through this lens looks something like what is shown in Figure 4.

	Same Place	Different Places
Same Time	**Audience-centric Crowdsourcing** *Audience-played games (e.g., pong, and flight simulator); Audience backchannels*	**Event-centric Crowdsourcing** *DARPA Red Balloons contest; Innovation Jams*
Different Times	**Geocentric Crowdsourcing** *Cyclopath; FixMyStreet; FourSquare*	**Global Crowdsourcing** *Wikipedia; ESP Game (AKA Google Image Labeler) and other Games With a Purpose*

Figure 4. Four quadrant model applied to crowdsourcing.

Global (any time, any place) Crowdsourcing

I've termed the lower right *different times – different places* quandrant **global crowdsourcing**, as this includes most commonly cited examples of crowdsourcing. In these cases, neither the spatial location and distribution of the crowd, nor the time at which members of the crowd become involved matters. Or to put it differently, there are no times or places from whence it is *not* appropriate to participate. Examples are Wikipedia (e.g., [3, 16]) the ESP Game [17], and other "games with a purpose" (GWAPs) [8]. It is clear that Wikipedia fits in this quadrant, but some explanation is needed for the ESP Game. While the ESP Game involves pairing up players to label images during a synchronous game – thus calling into question the fit of the "different times" dimension – the ESP Game's results are produced over multiple iterations of the game; furthermore, the design of the ESP Game makes clever use of bots to eliminate the necessity for synchronous human presence. The same may be said for other GWAPS.

Event-centric Crowdsourcing

Moving counterclockwise around Figure 4 brings us to the s*ame time – different places* quadrant: **event-centric crowdsourcing**. These are cases where a crowd is recruited for a particular event that has a start and a finish. One example is IBM's Innovation Jam offering, in which online events involving tens to hundreds of thousands of people lasting for several days are used to brainstorm on a set of topics (e.g., [2]) Another example is the DARPA Red Balloons contest [5] in which teams were invited to use social media to locate 10 red weather balloons displayed around the country for a prize of $40,000. The winning team devised their own contest, which split the prize among those who first reported a balloon, those who referred the first reporters, and those who referred the referrers of the first reporters [13]. In these cases, the crowdsourcing activity is organized around the event with its features, such as its start and end, driving the crowd's activities.

Audience-centric Crowdsourcing

Continuing around Figure 4 brings us to the *same place - same time* quadrant: **audience-centric crowdsourcing**. Crowdsourcing systems in this area are relatively undeveloped, but it represents a very interesting domain. One example is the genre of audience-played games sometimes found at technology oriented conferences; for example, Kelly [11] describes audiences divided into subgroups using individually held controllers to collectively play games like pong and control flight simulators. A different type of example, also found at technology-oriented conferences, is the use of chat or Twitter as a digital backchannel by the audience (for many examples see [4]).

Geocentric Crowdsourcing

A final counterclockwise move brings us to the *same place - different times* quadrant: **geocentric crowdsourcing**. Here the work of the crowd is focused on a particular place or geospatial region. The examples discussed in the previous section – Investigate Your MP, FixMyStreet, and Cyclopath – all fall into this quadrant. Another example of growing popularity is FourSquare [7], a GPS-enabled system that enables users to create and "check in" to establishments near their physical locale, and to provide information and tips which can be accessed by others. However Foursquare is more of a generic service (like Google Maps/cycling) than an application rooted in a particular place, and perhaps as a result uses game-like incentive mechanisms.

Summary

I view all but the lower right quandrant as examples of situated crowdsourcing, in which the crowd is associated with some context: a single place, a single event, or a single event in a single place. This situatedness is important because the context for the crowdsourcing offers a resource for structuring the activity of the crowd. Thus, in audience pong, the left half of the audience will play one side, the right half the other, and all will respond to what is visible in real time on the screen. In same time - different places situations, the temporal structure of the event shapes the crowd's activity. And in the geocentric case (different times - same place), the structure of the place itself can offer a way of organizing and focusing activity.

Computational vs. Deliberative Crowdsourcing

Another way in which I distinguish types of crowdsourcing systems has to do with the granularity of tasks the members of the crowd perform. In GWAPs, like the ESP Game, the users' tasks – e.g., generating textual labels suggested by an image – are very simple. In Wikipedia, the users' tasks are often quite complex – reorganizing an article, adding content, restructuring a paragraph. The degree of complexity of these tasks has consequences. One is that simple tasks are easy to compare and integrate. Thus, the ESP Game works by aggregating the answers for a given image: frequently given labels are likely to be high quality; labels generated only once are not. More generally, when the basic task in the crowdsourcing system is simple, digital mechanisms can be used to integrate the work of the crowd. In contrast, when tasks are more complex, as in Wikipedia, digital integration mechanisms rarely suffice. Integration and quality control occurs via social interaction, which can range from "edit wars" [16] to, more commonly, discussion and negotiation on the appropriate talk pages.

Another consequence of the nature of the users' tasks may be in the incentive systems necessary to motivate crowd activity. When the tasks are simple, as in The ESP Game, it may be that game-like incentive mechanisms are necessary. Even though an ESP Game player may recognize that generating textual labels for images is useful, the task is sufficiently simple that it does not seem very intrinsically motivating – it is only with the application of game mechanics that are unrelated to the larger endeavor that a crowd can be motivated to generate image labels. In contrast, more complex tasks are likely to be more intrinsically rewarding, and in addition, the social interaction required for integration and quality control is something that many find motivating in and of itself.

Closing Remarks

I am not sure if these are the best – or even satisfactory ways – of characterizing crowdsourcing systems. But I do believe the discussion is an interesting one, and it seems to me that there may a number of dimensions – user task complexity, incentive mechanism, result integration mechanism, how the higher level task is decomposed into more elementary subtasks, how the crowd is divided into cohorts, how cohorts are focused on particular subtasks – that the workshop could explore with an eye to better characterizing crowdsourcing systems.

REFERENCES

1. Baecker, R. M., Grudin, J., Buxton, W. A. S., Greenberg, S. 1995 *Readings in Human-Computer Interaction: Towards the Year 2000* (Second Edition) Morgan Kaufmann Publishers, Inc. See page 742.
2. Bjelland, O. M., and Wood, R. C. An Inside View of IBM's 'Innovation Jam.' *MIT Sloan Management Review*. October 1, 2008.
3. Bryant, S. L., Forte, A., and Bruckman, A. 2005. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. *Proc. GROUP 2005*. ACM Press (2005), 1-10.
4. Conference Backchannels. http://en.wikipedia.org/wiki/Backchannel
5. DARPA Red Balloons Contest. https://networkchallenge.darpa.mil/rules.aspx
6. FixMyStreet. Fixmystreet.org.uk
7. FourSquare. PlayFourSquare.com.
8. Games With a Purpose. http://www.gwap.com/gwap/
9. Investigate Your MP. http://mps-expenses.guardian.co.uk/
10. Johansen, R. 1988 *Groupware: Computer Support for Business Teams*. The Free Press.
11. Kelly, K. *Out of Control: The Rise of Neo-Biological Civilization*. Chapter 2. Addison-Wesley, 1994.
12. King, S. F. and Brown, P. Fix My Street or Else: Using the Internet to Voice Local Public Service Concerns. Proc. ICEGOV (2007). ACM Press (2007), 72-80.
13. MIT Red Balloon Challenge Team FAQ. http://balloon.media.mit.edu/mit/faq/
14. Panciera, K., Preidhorsky, R., Erickson, T., and Terveen, L. Lurking? Cyclopaths? A Quantitative Lifecycle Analysis of User Behavior in a Geowiki. *Proc. CHI 2010*. ACM Press (2010), 1917-1926.
15. Priedhorsky, R. and Terveen, L. The computational geowiki: What, why, and how. *Proc. CSCW 2008*. ACM Press (2008), 267–276.
16. Viégas, F. B., Wattenberg, M., and Dave, K. 2004. Studying cooperation and conflict between authors with *history flow* visualizations. *Proc. CHI* 2004. ACM Press (2004), 575-582.
17. von Ahn, L and Dabbish, L. 2004. Labeling images with a computer game. *Proc. CHI 2004.* ACM Press (2004), 319-326.