Collective Intelligence for the Design of Emergency Response

Adriana S. Vivacqua
PPGI/DCC-IM – Graduate Program in Informatics
Department of Computer Science, Institute of Mathematics
UFRJ – Federal University of Rio de Janeiro
Rio de Janeiro, Brazil
avivacqua@dcc.ufrj.br

Marcos R.S. Borges
PPGI/DCC-IM – Graduate Program in Informatics
Department of Computer Science, Institute of Mathematics
NCE – Electronic Computing Institute
UFRJ – Federal University of Rio de Janeiro
Rio de Janeiro, Brazil
mborges@nce.ufrj.br

Abstract— **Collective Intelligence has received much attention in recent years, as organizations and businesses discover the power of crowds. With new technologies, such as blogs, twitter, wikis, photo sharing, collaborative tagging and social networking sites, people may create and disseminate content as never before. Many organizations are looking for ways to harness this power, which is being called *collective intelligence*. Additionally, many open-source initiatives have shown that it is possible to obtain high quality results from collectively produced work. In this paper, we consider the domain of emergency response. We have been working in this domain for a few years now, mostly focusing on small disaster relief groups and field command. One of the command problems is lack of information about the location where the disaster occurred. We believe some of this information could be *crowdsourced*, changing current practices and enabling more effective response to the situation. In this paper, we discuss the possibilities for *crowdsourcing* and present an initial proposal for harnessing collective intelligence for the fire department.**

Keywords-collective intelligence; crowdsourcing; disaster relief

I. INTRODUCTION

Emergencies and disasters happen all over the globe. Floods, hurricanes, fires and landslides threaten people's lives and must be dealt with in a speedy fashion. Emergency relief teams are heavily trained to deal with adverse situations, but they need information about the environment to be able to make appropriate plans and act to resolve the situation.

On a different end of the spectrum, large numbers of people are now online. With the advances in networking technology and the dissemination of the internet, a culture of content production and sharing is being created. Many people write post on blogs, wikis or twitter and have an active profile on online social networking sites such as Facebook[1] or Orkut[2]. The current generation of internet users has gotten used to online participation in many different forms.

The participation of large groups of people is sometimes able to generate results above and beyond what a single individual could accomplish [13]. One example is Wikipedia, which is completely generated by volunteers and Linux, one of the best known open source systems aver written. Harnessing this power to one's advantage is something many organizations have been trying to do. One aspect of collective intelligence that has gotten much attention recently is *crowdsourcing*.

Crowdsourcing (outsourcing to a crowd) is giving a task for a crowd to execute, instead of executing it oneself. For instance, developing alternate solutions for a given problem could be done by a sole expert or could, instead, be conducted by a large set of individuals, given appropriate tools for aggregation of the different contributions. This works especially well in situations where opinions are needed, and can be obtained by filling a questionnaire, for instance. Amazon recently launched a service called the Mechanical Turk[3], a "marketplace for work", where it is possible to *crowdsource* any type of task. Individuals post "task offers" for anyone in the community to perform, for a small payment (tasks include testing systems, proofing websites, etc.). This is currently being explored by the HCI community, who is trying to determine if it is an effective platform for running evaluations (for example, by providing an application and a link to a usability survey to be filled for a small payment.)

Technology is also becoming easier to use and APIs allow access to predefined content and the creation of "*mashups*", or sites that combine content from different sources. This enables processing and reuse of different data beyond what it was originally intended to. Google Maps is one example of an application whit an API that has been used in a number of *mashup* applications very successfully.

Our group has a history of working with the Fire Department to develop technology to support them. One of the problems they face has to do with the lack of information about the area surrounding the location of the emergency. We feel

[1] http://www.facebook.com/
[2] http://www.orkut.com
[3] https://www.mturk.com/mturk/welcome

978-1-4244-6763-1/10/$26.00 ©2010 IEEE

this problem could be tackled through collective intelligence, and have initiated a project to address this possibility.

Our goal is to investigate in which ways collective intelligence could be tapped for emergency management. In this paper, we present an analysis of this potential usage. In the next section, we present the activities involved in disaster management and what types of problems exist. In Section 3 we discuss Collective Intelligence and Crowdsourcing, and present some initiatives that bring disaster management and collective intelligence together. We follow by presenting our take on how collective intelligence could help disaster relief efforts, and present a scenario based on a real problem. We make some final considerations on the subject and outline future directions for work.

II. Disaster Relief

Aligne [1] defines a crisis as a situation in which there is a break from previous events. A crisis threatens the functioning and values of an individual or community, and there is an urgent need for action despite difficulties. One of the most important characteristics of a crisis is that it is unpredictable and cannot be totally anticipated by scenarios. To resolve a crisis, it is important to reflect and gather information with which to mobilize resources. Crisis management is modeled as a cycle in which several activities take place. These activities happen in three moments: pre-disaster, response (during the disaster) and post-disaster [1]. The cycle and activities are shown in Figure 1.

Figure 1. Cristis Management Activities (from [1])

Emergency response usually involves multiple agencies, so coordination becomes a critical aspect of the relief effort [6]. This gets very complicated due to the very different information systems used by these organizations [2]. Furthermore, typically all normal networks and power sources are unavailable at the site of the disaster.

An outline of an arrangement of resources is presented in [7] (Figure 2). Each Organization involved in a relief mission has its own communication network. Communication between people is mediated by each organizations' centralized nodes (i.e., there is no direct communication between a firefighter and a policeman in the field). Firemen, Police and Medical organizations have people in the field, whereas specialists are remotely consulted and the Government has national support networks.

In the lower level, difficulties are great: wired or cellular networks cannot be expected to work at the disaster site, and electrical power may not be available either. Researchers have explored solutions using Mobile Ad-hoc NETworks (MANETs), in which connectivity among devices is unreliable and changing: the distance between devices cannot be over a few hundred meters in open spaces. However, MANETs do not require any communication infrastructure and they can be deployed on portable devices operating with batteries, such as Tablet PCs and PDAs. Many problems are still open with such networks, e.g., routing protocols, extending operational use time of the devices, dynamic discovery of nodes, and security.

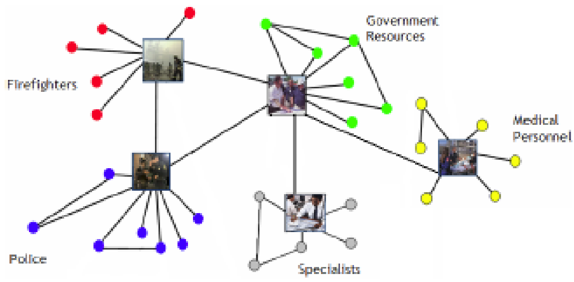

Figure 2. Cristis Management Organizations (from [7])

Recent studies have noted that people tend to use their social networks to communicate and keep each other informed in the event of a large scale emergency [14]. Individuals use twitter, SMS, cell phones and the web to make contact and keep each other abreast of the situation. It has also been noted that groups present convergent behavior in times of emergency [4]. It makes sense, then to tap into this collective intelligence during disaster situations.

III. Collective Inteligence and Crowdsourcing

In recent years, researchers have proposed that large groups can perform as well or better than individual experts. Surowiecky [13] argued that groups with certain characteristics are almost certain to produce useful results, and named it the Wisdom of Crowds. One example he provides is the stock market (the futures market in particular), where traders buy and sell stock according to what they believe the future holds for a company, given the information they have. Prices in these markets are not set by one organization or through a coordinated effort, but by thousands of individuals acting upon their information to achieve what they believe will be the best outcome for themselves [9]. Futures markets have been shown to be better predictors of future events than expert opinions in some cases [13]. Based on this observation, some approaches have been created that use this market approach to elicit information and make predictions, these are usually called information or prediction markets. Other authors note that, with increased online participation, the amount of data available has increased greatly, and it is possible to make use of this data through the application of artificial intelligence methods [9].

Another well known case is the open source initiative. In open source development, several people contribute to one joint effort, frequently generating a solid product. Linux is probably the best known, but there are many others. Wikipedia is another good example: it is a collectively created encyclopedia. After the Wikipedia's success, wikis have been started in several subdomains (travel, dictionaries, books, etc.), proving that the approach is valid in many cases. This community construction approach was also used to construct a common sense knowledge based at MIT, which reached in a few months the size of another well known common sense knowledge base, CYC, built the traditional way [11].

Collective intelligence involves combining knowledge (which may include behaviors, preferences, ideas, etc.) from a group of people to produce novel information or insight [9]. Sunstein presents four ways in which groups could elicit the information they need [12]: through statistical averages of independent contributions; through deliberation and reasoned exchange of facts and ideas, to improve on the individual judgments; using a pricing system or a market; or via the internet, to elicit information from whoever wants to contribute. The internet opens up new possibilities for interaction and knowledge construction by large groups of people. This knowledge can be put to use in many situations, including disaster relief.

A. *Public Participation in Emergency Magnagement*

Some projects have been initiated to gather information about a certain physical environment, to help (or pressure) authorities resolve the situation. One of the precursors was the *Better Pune*, an initiative of the Better Roads Group, an NGO in India to help local authorities take action. The system allows citizens to post on a Google Maps interface potholes and drain clogging they encounter in the city[4]. Similar initiatives have appeared up in other locations. *FillThatHole*[5] is a UK based site sponsored by the UKs's National Cyclists Association. Users report report potholes using a Google Maps mashup, and this information is sent to the responsible authorities for action, generating an official report. On the map, potholes marks are color coded according to their current status, when they are fixed, authorities send this information back to the CTC and the map is updated.

Another well structured initiative is *FireMash*[6] (shown in Figure 3). This is a system to report bushfires in Australia, a frequently devastating event. Their point is that, through early warning, appropriate steps can be taken to deal with the situation (e.g., evacuation) before tragedy strikes. FireMash combines live feeds from the New South Wales (NSW) Rural Fire Service (they currently provide information on both twitter and RSS) and maps their location on a Google Map. Additionally, any citizen may tweet a fire using the hashtag #nswfires and the system will also post it to the map. Perhaps the most useful feature, people may inform the location of their houses and the system will inform them when a fire is getting close to that location.

Earlier this year, Google, Microsoft, Yahoo!, NASA and the WorldBank sponsored an event where hackers gathered to generate systems that could be useful for people facing emergency situations [8]. During this event, many ideas for citizen empowerment in crisis situations were generated, and a few were implemented in a two day workshop.

Figure 3. FireMash, an application that allows the population to report wildfires in Australia via Twitter.

[4] http://www.punecorporation.org/betterroads/Home/Home.aspx
[5] http://fillthathole.org.uk/
[6] http://www.firemash.com/

Liu and Palen present a survey of current technology based on spatiotemporal mashups to support crisis management [5]. They provide an analysis of systems that pinpoint occurrences of a wide range of events, including domestic violence, earthquakes, terrorist attacks, school shootings and the rising sea level. They conclude that visualizations of spatiotemporal data are valuable assets for emergency management.

IV. Collective Intelligence for Disaster Relief

One of the challenges faced by disaster relief organizations is gathering information upon which to act. Many such organizations do not have the manpower to conduct extensive research in order to be better prepared at the moment of a disaster. Additionally, large cities often grow without any order, which makes it difficult for organizations to keep track of changes in the cityscape. We believe this need could be addressed by a crowd.

In this section we discuss different approaches to collective intelligence for disaster relief. We consider the information needs and possibilities for collective participation in five phases of the emergency management activity:

- Prevention (pre-disaster): involves attempting to avoid the emergency altogether. To avoid an emergency, responders need information about the chances of certain hazards turning into emergencies. This means verifying the environment to check whether there are any high-risk conditions that could be dealt with to avoid a disaster. One example would be identifying unsafe wiring in a power plant before it short circuits or hurts someone. In this phase, there is also a need for risk assessment, so a team leader may judge how urgent the situation really is.
- Mitigation (pre-disaster): involves minimizing the effects of a potential disaster. During this phase, an emergency management team must have information about the effects of a given emergency to attempt to minimize these effects if it were to happen. This entails assessing the risks and potential consequences of an emergency. In this phase, knowledge about the area surrounding a potential emergency is necessary in order to take appropriate action. This also involves keeping abreast of the current state of the environment, which may be subject to frequent change. For instance, should new housing be constructed in the surrounding area of an oil plant, an accident, such as an oil spill, could potentially affect the people who live there.
- Preparedness (pre-disaster): involves planning the best way to respond to a disaster, before it strikes. It entails being ready to deal with a situation. This usually leads to personnel training, designing emergency handling procedures or plans and running scenarios to simulate emergency situations. Information needed at this stage has to do with the area surrounding the emergency, and what resources are available to deal with it (e.g., where are the fire extinguishers located?)
- Response (disaster): the actual efforts to minimize the hazards created by the disaster. During response, the emergency team is dealing with the situation as it evolves. It should be noted that actions reflect on the situation, so the context is constantly changing. It is important to keep up with these changes in the environment, in order to take appropriate action, but it is hard for a team to keep the environment under watch at all times. Information such as a rooftop collapsing in a section that wasn't being closely monitored, or someone being trapped by debris at a certain location is necessary for respondents to decide on which action to take.
- Recovery (post-disaster): recovery involves dealing with the fallout and returning the situation to normal. This collecting information about the area and surroundings to see what isn't functioning as it should and dispatching services to these locations to normalize the situation. Having a comprehensive picture might also allow respondents to maximize their resources, dealing with similar or related problems as a group instead of dealing with each one individually.

As noted, all phases of crisis management require a lot of information for appropriate decision making. We believe that, by focusing the efforts of the population, it is possible to generate useful, actionable information that respondents can use. Some examples have been mentioned in the previous section.

To provide opportunity for ample participation, systems would need to be setup to enable the population to provide information either via the web or mobile phone. We now discuss possibilities for citizen participation in emergency management in light of the four approaches proposed by Sunstein [12], fitting the activities into each approach:

- Statistical – statistical approaches attempt to combine several contributions in a mathematical way (for instance, averaging values provided). Voting systems can take the information provided by the community and aggregate it. These systems can be used both in pre-disaster and in post-disaster moments. People would be allowed to decide on a number of given options (e.g., which area to restore first). Another possibility for aggregation would be averaging a number of GPS coordinates sent to a central server to mark certain elements (e.g., potholes): if multiple senders mark the same element, its position can be calculated as a function of the coordinates sent in.
- Deliberative – deliberation involves discussion and exchange of ideas, which should improve on individual judgments. These could be easily implemented as discussion forums that follow on to voting systems. These systems, however, are

slow to yield results, as people take time to deliberate and finally vote. These types of systems could be better used in pre-disaster stages, for prevention or mitigation, when participants have more time to discuss.

- Market – markets are in a sense, similar to betting games: participants place their bets on the options they believe will come true given the information they have. The market regulates itself, as individuals buy or sell "stock" of each option. These work well when there are definite answers and when these are verifiable at a later date. Markets could be used, for instance, to predict which areas will be most affected when a storm hits. This could assist in the prevention of disasters. This could also assist during response, if individuals evaluate which area of the current emergency will experience certain consequences first (e.g. roof collapse) or will have the worst consequence (e.g. someone will be trapped). This would enable individuals (onlookers) with different points of view to contribute, leading to better coverage of the affected area.
- Volunteer contributions – through the internet or via mobile phones, individuals may contribute information as they see fit. Focusing these contributions (as, for instance, in the FireMash example presented before) is the challenge emergency managers must address to turn this into useful information. Using twitter, SMS or the web enables any citizen to provide information. This information could be used at any stage of the emergency response, and these methods usually generate a large amount of information. Onlookers could, for instance, map the environment (both resources and imminent threats) and this information could be sent to respondents as they travel to the disaster location, and arrive there with better knowledge of what expects them. Sending photos or videos would greatly help the authorities better understand the situation, without having to rely solely on individual judgment.

These approaches are bound to generate a large amount of information (the last one in particular), which risks worsening respondents' decision-making process. There is a possibility of information overload. One way to address this is to apply data mining and information retrieval techniques to cluster and pre-process this information before presenting it to respondents, as suggested in [10].

A second problem is the reliability of the incoming information. It is important to provide correct information so as not to waste respondents' efforts. Given that information is sent in, and not verified in any way, there is no way to know if the information is true. However, if many individuals send in the same information, this can be seen as corroboration, and that information can be considered more reliable.

Regardless of the approach adopted, an information central would be necessary, so that incoming information may be analyzed and retransmitted to the authorities.

A. Scenario: the Hydrant Map

When an emergency is reported, response teams are dispatched to the location and make an assessment of the surrounding environment and resources available to deal with the situation. In many cases, the responding agency does not have information about what it will encounter at the scene of the emergency (e.g.: Are there hydrants? Where are they located? Where are the fire extinguishers?) This information is crucial for the relief mission to function well, but the agencies lack the resources and manpower to collect it in advance. Frequently, upon arrival to a disaster site, the first action a commander takes is to send a scout to locate all fire hydrants in the neighborhood, because he or she doesn't have sufficient knowledge about the environment to make appropriate decisions. In this section, we provide an example of collective intelligence applied to emergency management, using this scenario.

One solution to this problem could be to *crowdsource* the task of locating hydrants, creating a comprehensive database that the Fire Department could rely on when responding to a fire. This means offering a service through which the population at large can provide information about the location of fire hydrants in their vicinity. This means designing an information collection interface to enable popular participation on this task.

A web based *mashup* interface would work in a similar way as the ones described in previous sections. Individuals could enter information about the hydrants they knew of and their locations by dropping them on a map interface. Any additional information could be provided through a simple form.

A more interesting interface would be to have a service allowing the population to send in messages reporting a "hydrant sighting", This message would contain their location (automatically captured from GPS-enabled phones) and the fact that a hydrant was present. The visualization would follow the standard Google Maps *mashup*, as shown in Figure 4, and would be the same for both collection interfaces.

On the other end of this application, useful information needs to be available to firefighters when necessary. The information collected could be saved in a database, for querying when an emergency arose.

However, this information would also need to be aggregated. We would expect to receive several reports of the same hydrant, and that these would form clusters around certain locations, roughly defining the correct position. The location could be calculated based on the number of messages coming from each vicinity, or a general area could be provided, within which the hydrant is located. This, of course, cannot be too large, or the application loses its usefulness. It is important that, upon receipt of an emergency, the response team be able to get there and find the resources where they are, speeding up the response process.

Figure 4. Hydrant Map visualization

V. FINAL CONSIDERATIONS

As seen in recent years, disasters and emergencies generate a lot of spontaneous participation [4][14]. Individuals tend to communicate with others about what is going on, and, given the increase in citizen journalism seen in recent years, this tendency is bound to increase as well. People use twitter, blogs and wikis to report what they have experienced, sometimes making photos and video as well. The challenge is creating synergy and directing this energy towards the creation of content that could be useful for emergency management.

In both pre and post-disaster stages, many efforts are possible, from reporting of local resources to assessment of the status of the location. During response, the opportunities are fewer, but still exist. Respondents have stated that they rely on the population for information when they arrive at a site. This helps them assess the situation and decide where to act (for instance, knowing that someone is trapped and where after a landslide is important to guide the operation.) This information could be provided by multiple people in multiple locations, and aggregated to help pool resources and deal with the situation more efficiently.

Crisis management requires a lot of information for appropriate decision making. However, researchers have also noted that crisis situations tend to generate a great deal of information, which, in turn, frequently leads to information overload by decision makers [3]. This means that decision makers receive more information than they can possibly process, especially if considering the time constraints imposed by an emergency situation. The point is to have information at the right time and place, turning it into an actionable resource. In many situations, the unconstrained generation of information by the population may contribute to this information overflow.

Surface mapping, such as the scenario presented, is a real problem in a large city such as Rio de Janeiro, where authorities do not have the resources to resolve it, yet it is a problem easily resolved by a mass of people who routinely walk the city streets. The population could be enlisted to help in this and other efforts. In this paper, we present an initial mapping of how to adopt collective intelligence techniques to assist crisis management. Our group has been exploring emergency management for a while now, but there is still much to be done. Our next steps will be to map, for a given domain, the information necessary for operational decision making, and to construct applications such as the one described to verify their feasibility. We believe that this is an approach that may generate fruitful results.

ACKNOWLEDGEMENTS

Marcos R.S. Borges was partially supported by grants No. 304252/2008-5 from CNPq (Brazil) and No. R0308LAC004 from LACCIR.

REFERENCES

[1] Aligne, F. Which Information and Decision Support System for the Crisis Management? Information Systems and Technology Panel (IST) Symposium. Bucharest, Romania, May 2009

[2] Carminatti, N., Borges, M. R. S., Gomes, J.O.: "Analyzing Approaches to Collective Knowledge Recall". Computer and Informatics 25(6), 2006, 547-570.

[3] Hoare, G., Nield, J., Belcuore, T., Rich, T. Information Needs and Decision Support in Health and Medical Disasters. In Proceedings of the 5th International ISCRAM Conference – Washington, DC, USA, May 2008.

[4] Hughes, A.L., Palen, L., Sutton, J., Liu, S.B., Vieweg, S. "Site-Seeing" in Disaster: An Examination of On-Line Social Convergence In Proceedings of the 5th International ISCRAM Conference – Washington, DC, USA, May 2008

[5] Liu, S., Palen, L. Spatiotemporal Mashups: A Survey of Current Tools to Inform Next Generation Crisis Support. In Proceedings of Conference on Information Systems for Crisis Response and Management (ISCRAM) 2009

[6] Llavador, M., Letelier, P., Penadés, M. C., Canós, J. H., Borges, M. R. S., Solis, C.: "Precise yet Flexible Specification of Emergency Resolution Procedures". 3rd International Information Systems for Crisis Response and Management Conference, New Jersey, USA, May 2006, 110-120 (Best Article Award).

[7] Ochoa, S.F., Neyem, A., Pino, J.A., Borges, M.R.: "Supporting Group Decision Making and Coordination in Urban Disasters Relief Efforts". Journal of Decision Systems 16(2), 2007, 143-172.

[8] Random Hacks of Kindness. Website available at : https://sites.google.com/a/rhok.org/www/home

[9] Segaran, T. Programming Collective Intelligence: Building Smart Web 2.0 Applications. O'Reilly, 2007

[10] Silva, C. R. C. ; Tavares, T. C. ; Garcia, A. C. B. ; Nogueira, J. L. . Governo Eletrônico em Ambientes Colaborativos Virtuais. In: SBSI 2009 V Simpósio Brasileiro de Sistemas de Informação - Workshop de Teses e Dissertações, 2009, Brasilia. Anais do SBSI 2009 - V Simpósio Brasileiro de Sistemas de Informação, 2009.

[11] Singh, P. The Public Acquisition of Commonsense Knowledge. Proc. 2002 AAAI Spring Symposium on Acquiring (and Using) Linguistic (and World) Knowledge for Information Access, 47-52.

[12] Sunstein, C. Infotopia: How many minds produce knowledge. Oxford University Press, 2006

[13] Surowiecky, J. The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations, Random House, 2004

[14] Vieweg, S., Palen, L., Liu, S.B., Hughes, A.L., Sutton, J. Collective Intelligence in Disaster: Examination of the Phenomenon in the Aftermath of the 2007 Virginia Tech Shooting. In Proceedings of the 5th International ISCRAM Conference – Washington, DC, USA, May 2008