Global Manhunt Pushes the Limits of Social Mobilization

Iyad Rahwan,[1,2] Sohan Dsouza,[1] Alex Rutherford,[1]
Victor Naroditskiy,[3] James McInerney,[3] Matteo Venanzi,[3]
Nicholas R. Jennings,[3] Manuel Cebrian[4,5]

[1]Masdar Institute of Science & Technology, UAE
[2]University of Edinburgh, UK
[3]University of Southampton, UK
[4]NICTA, Melbourne, Australia
[5]Massachusetts Institute of Technology, USA

Abstract

The Tag Challenge, sponsored by the US State Department, required locating and photographing 5 target persons in 5 different cities in the United States and Europe in less than 12 hours, using only their mug shots. We report on how we used social media to win this challenge, finding 3 of the targets. We reflect on the difficulties we faced, the lessons we learned, and the implications of this capability.

1 Introduction

Back in 2009, to commemorate the 40th anniversary of the creation of the Internet, the US Defense Advanced Research Project Agency (DARPA) launched its Network Challenge (also known as the "Red Balloon Challenge"). The challenge aimed to test the limits of social mobilization and rapid information gathering using social media. It required competitors to locate 10 weather balloons tethered at random locations all over the US. The winning team, based at MIT, found all balloons within 9 hours by mobilization through social media [18]. The winning strategy relied on a novel incentive scheme in which people were rewarded both for reporting balloon sightings, as well as for recruiting their friends to look for balloons [17]. Further theoretical work proved that the strategy is in fact optimal in terms of minimizing the investment to recover the information [4], and is also the most robust to misinformation [16].

In March 2012, the "Tag Challenge," funded by the US State Department, raised the bar significantly higher (`www.tag-challenge.com`). The challenge set the unprecedented task of locating and photographing 5 people in different cities, across two continents, within 12 hours. There was a reward of $5000 for the winning team. Only a single mugshot of each target person (or "suspect") was released to the teams at 8am local time on the day of the competition (Figure 1). Each volunteer target wore a t-shirt with the competition logo (the appearance of which was also not known until the first mugshot was released) and was instructed to follow a 12-hour itinerary designed to reflect a normal day. For example, the New York City suspect started at Columbia University, had breakfast at a cafe nearby, took the subway to the World Trade Center site, then went shopping, and so on. As such, the task represented a realistic search for an individual following a characteristic mobility pattern. The other suspects were in Washington D.C., London, Stockholm, and Bratislava.

Figure 1: Only information given was a single mug shot per city released at 8:00am local time

Now, the Tag Challenge is significantly harder than the Red Balloon Challenge because it required locating people in extremely populated cities, where the pace of life can help people "hide in plain sight" [15]. Therefore, the target people were much harder to spot than large red balloons. Furthermore, people are mobile, making it difficult to rule out locations that have already been visited. While the suspects were not explicitly hiding or in disguise, searching for one moving about in a city like New York, with a population of 8.2 million people, seemed nearly impossible. In particular, it was an open question as to whether it is possible to use social media to accomplish such a difficult task, as evidenced by the difficulty of finding suspects in police investigations [19].

Despite these challenges, our team won the challenge by locating 3 of the targets using a Web platform, a mobile application, and an incentive scheme. This was accomplished without any of the team members being located in any of the target cities. Specifically, Figure 2 shows the targets our team found, and the approximate local time at which we submitted their photos to the organizers. The targets in London and Stockholm remained at large, although pursuing them after the allotted 12 hours was not part of the competition.

Other teams used Web sites and social media in combination with a wide range of approaches, including attention raising measures (e.g., search engine optimization + broadcasting a high volume of messages on social media), standard incentive measures (e.g., promising shares of revenue to participants + pledging winnings to charity), and deceptive measures (e.g., impersonating the challenge organizers + attempting to sabotage rivals teams with a virus). Clearly, there was a wide range of possible strategies in this competition; our approach was based on an understanding of the key challenges of rapid mobilisation of crowdsourcing teams. In the next section, we give an overview of these challenges.

2 Challenges

There are at least three distinct problems in crowdsourcing rapid information gathering: (1) the *mobilization* of participants; (2) the *aggregation* of information; (3) and the *verification* of information. We discuss each one in turn.

2.1 Mobilization

The success of search in social mobilization clearly requires individuals to be motivated to conduct the search, and to participate in the information diffusion. In an attempt to replicate Stanley Milgram's "small world experiment" [14], it was observed that the majority of message forwarding chains observed empirically terminate prematurely. Specifically, Dodds et al. conclude that "*although global social networks are, in principle, searchable, actual success depends sensitively on individual incentives*" [5]. In other words, a key challenge in social mobilization is the *incentive challenge*. It has also been observed that the success of crowdsourcing mechanisms, in general, can vary depending on the details of the financial incentive scheme in place [12].

Figure 2: Location of the 5 targets, and local time at which the 3 targets were found

While mobilization requires people to be motivated, in natural disaster situations there seems to be a very strong intrinsic desire to contribute. Once people identify a focal point for submitting information (e.g., Ushahidi[1] or Cosm[2]), it becomes a hub and everyone contributes and invites other contributors. So motivation for participation and spreading information is less of a concern in such scenarios, and the efforts can be focused on information aggregation and verification.

However, there are many situations in which the level of intrinsic motivation is not as strong or widespread. Consider the problem of finding a single missing person or criminal (as in the Tag Challenge). The sparsity of the information being sought suggests the need to motivate people to *route* the problem to others who are better placed to report on it (e.g., to someone who lives in the same area as the target or who is likely to come across them in the course of daily life). This is a much harder task, because the person who eventually finds the target is only one element in a chain of participants. This is in contrast to more conventional crowdsourcing settings like Ushahidi, where the bulk of each task begins and ends with a single individual. In short, tasks on Ushahidi are independent, whereas each successful *search* task requires an uninterrupted chain of motivated participants.

The traditional method of finding missing persons or criminals requires a central actor to do all the work in advertising a reward or ransom (respectively), since people have no incentive to route the information to others. In fact, if someone spots the criminal or has useful information of their whereabouts, they even have incentive to hide that information from others to avoid sharing the reward. These features act as barriers to large-scale mobilization in situations where time and resources are critical. Even with the right incentives to motivate people to act, there are still open questions about the economics and efficiency of the crowdsourcing system. Much academic work in crowdsourcing is now exploring this problem, for example to optimize the efficiency of micro-task markets like Amazon Mechanical Turk or CrowdFlower [12].

[1] `ushahidi.com`
[2] `cosm.com`

2.2 Aggregation

Given a sufficient quantity of motivated participants, the next challenge is to aggregate their responses. Even if we assume, for now, that all contributions are reliable, visualizing and synthesizing the information into a form that is actionable requires careful thought.

This is the main challenge that platforms such as Ushahidi and Cosm seek to address [8]. If there is a forest fire or an earthquake, one can use Ushahidi to aggregate geo-tagged reports posted by thousands of people, then visualize what is going on using the resulting crisis map. The system has proven its efficiency in rapid information aggregation through various deployments, including the aftermath of the devastating earthquake in Haiti.

However, there are scenarios in which aggregation is less trivial. For example, when the information submitted has high volume, it is often necessary to identify which information is most relevant to the task at hand, in order to avoid clutter in the visualized maps. Things get even harder when aggregating different types of data, for example when combining spatial and temporal communication data to characterize response to massive emergencies [2].

2.3 Verification

Experience with Ushahidi, however, also highlighted the verification problem in crowdsourcing. Verification is important as people stand to benefit from submitting false reports, for instance to get more food or receive help sooner, often at the expense of others. In the case of search, false positives from well-intentioned participants, and malicious reports from those seeking rewards are also likely. The Metropolitan Police in London have recently released the photographs of 5,000 suspects believed to be involved in the London riots of 2011 [19]. Cases of mistaken identity are inevitable, especially matching against poor quality CCTV images. This is known to happen with general crime appeals from television programs such as Crimewatch in the UK, and America's Most Wanted in the USA.

In one of the first serious attempts to tackle automatic verification, the Ushahidi team developed a set of algorithms called SwiftRiver, that use machine learning techniques to classify information sources and content (e.g., Twitter accounts or individual tweets) to filter the important/relevant ones. Nevertheless, the problem of verification remains a very significant challenge, especially in contentious or competitive domains [13].

3 Our Approach

We formulated a strategy to address these challenges, focusing primarily on the mobilization aspect, as we believed the hardest part would be to get people to take action. We detail this strategy, and our approach to implementing it.

3.1 Strategy

Following the Red Balloon Challenge winning strategy, we used an incentive scheme that is designed to encourage two things simultaneously: (i) reporting to us if you found a target; (ii) helping recruit other people to search for the target. We described the strategy as follows: "If we win, you will receive $500 if you upload an image of a suspect that is accepted by the challenge organizers. If a friend you invited using your individualized referral link uploads an acceptable image of a suspect, you also get $100. Furthermore, recruiters get $1 for each recruit they refer to sign up with us through their individualized referral link, up to the first 2,000 recruits." The scheme is illustrated in Figure 3.

Here, the incentive to refer others had to be significant, because otherwise, attracting additional participants to the challenge would only decrease an individual's chances of

getting a reward. Thus, by paying people for referring their friends, as well as for finding suspects, the incentivized behavior changes fundamentally.

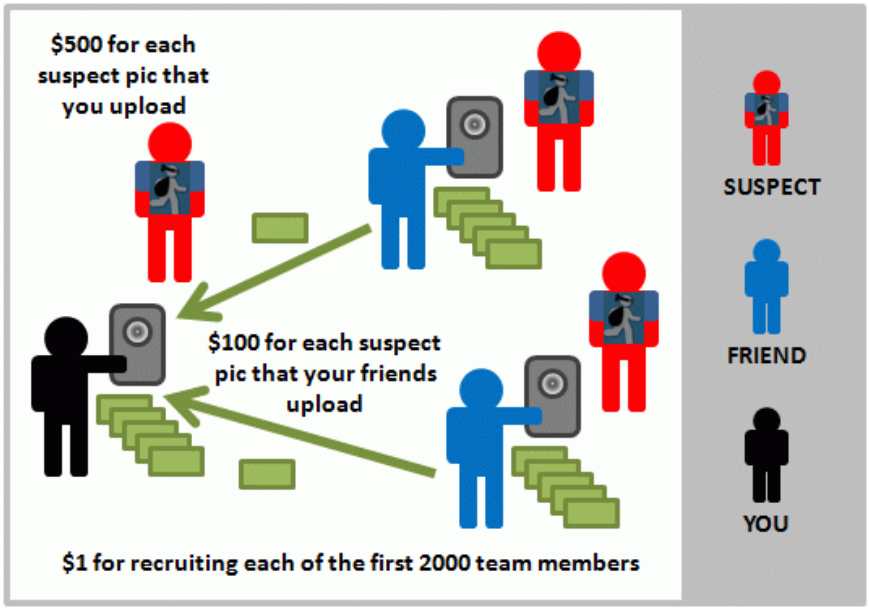

Figure 3: Description of the incentive scheme used in our strategy

We found that aggregation and verification were not major issues in this competition because the number of submissions we received was relatively small. In particular, we manually verified our submissions in the Tag Challenge, making sure the photograph of the target matched one of the mugshots provided earlier by the organizers (and the human brain is very effective at this task). There was a single submission from each city, so aggregation was not required.

The t-shirts of each target suspect had a unique and previously unknown code, giving the organizers the ability to tell whether the suspect had really been found. On our part, we had to make sure this code was legible from the picture. In every case, we contacted the submitters directly; if part of their backstory was inconsistent with expectations then we would not have trusted their submission. These activities were all possible because, again, the number of submissions was small. Clearly, such a detailed approach would not be possible at larger scale, although we have recently proposed a mechanism for crowdsourcing the verification tasks simultaneously with the search tasks [16].

3.2 Implementation

To facilitate our strategy, we designed and built a web application that allows people to submit photographs of suspects online and recruit other people into our campaign. To reward participants for recruitment, it was important to be able to trace a new recruit to the existing recruit who referred them (where applicable). This was accomplished by providing an individualized referral link to share on social networks, a process made easier through the site's template email and social media sharing links.

Figure 4 shows the distribution of clicks to our team's web site over the week leading up to the challenge. The figure clearly highlights that the areas of interest are generally focused around the target cities (New York, Washington, Stockholm, London and Bratislava).

We also built a mobile phone app that allows people to view the list of suspects as they are released, and submit pictures of them directly from their phones. The app was released on the Android Market. Social media platforms were not ignored, as we went

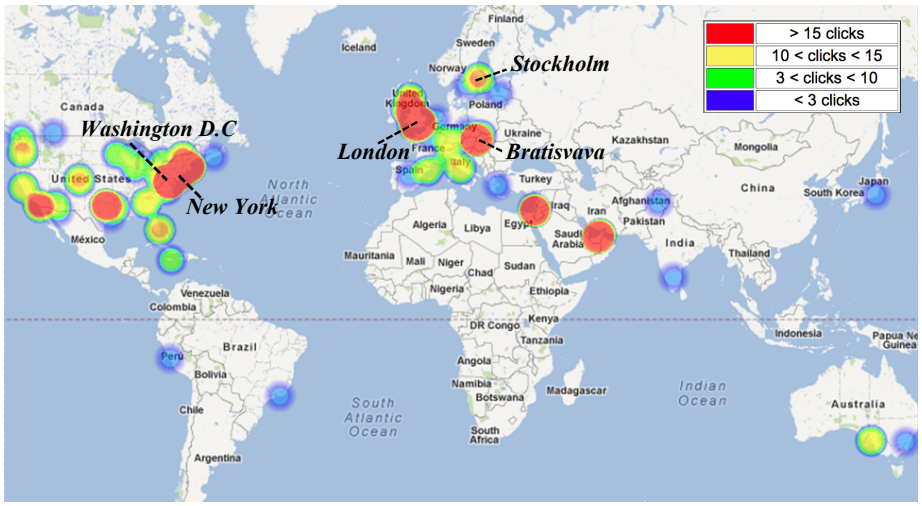

Figure 4: Heatmap showing distribution of visitors to our team's site

about creating a Twitter feed and Facebook page for people to *follow* and *like*, respectively, so they could receive updates.

Somewhat to our surprise, neither the web platform nor the mobile app were used very much for submission during the competition. Successful participants in our team preferred to submit their photos via email, suggesting that when important information needs to be reported, people value direct communication rather than going through official channels. Perhaps this could be explained on grounds of convenience (i.e., each new platform requires a learning curve for participants), or the lack of trust that highly valuable information (i.e., the winning photos) will not get lost in a sea of noise. The web site and the app are both components of a crowdsourcing platform being developed by our team, so addressing these issues are important for scalability.

Nevertheless, both the web platform and mobile app contributed significantly to our credibility by demonstrating commitment and Web presence. This helped significantly with the final element of our approach, which was to generate as much attention as possible for our team on blogs and news Web sites. We sent out press releases explaining our interest in the challenge, and posted notes about our team in related forums. We next explain how our strategy gave us an advantage over the other teams.

4 Other Teams

We were able to succeed by leveraging a combination of social media and information hubs, and by building up a reputation as a credible, reliable team. In this section, we highlight some of the tactics used by other teams, including deceptive ones, and reflect on their effectiveness and implications. Neither we nor any other team were able to find the suspects in London or Stockholm (despite significant coverage in the former).

We are aware of the existence of five serious opposing teams. Some teams (e.g., team `@WeTagChallenge`) used search engine optimization tactics to raise the profile of their team websites above that of other teams. On the other hand, another team with Twitter handle `@TAG_challenge` apparently attempted a "man in the middle" attack using a name that was very similar to the official `@TAGchallenge` account. We suspect that they may have been attempting to pose as the actual TAG Challenge organizers in order to intercept submissions, even to the extent of duplicating the official tweets.

Yet another team, with the Twitter handle `@TeamRave`, mostly ran its campaign on Twitter and its Web site. They promised the entire $1,000 share for each accepted suspect

picture, leaving no incentive for people to recruit others.

The most serious (and aggressive) competitor was "Tag Team." They used an incentive scheme similar to ours ($400 to finder, $100 to referrer) with an added twist – instead of using the remainder to build up a critical mass of recruits, they promised that the remaining $500 per suspect would go to charity. Their strategy for spreading awareness consisted predominantly of their Twitter account `@TagTeam_` surfing trending hashtags, and tweet-spamming many individuals, social, governmental and private organizations in the target cities, often with an explicit plea for a retweet. The vast majority of these were ignored and, we believe, reduced their credibility. Interestingly, they sent many tweets that mentioned the `@TAG_challenge` handle, which was, as described earlier, the handle of an imposter account.

In the days leading up to the challenge, Tag Team also tweeted messages targeted at our team, including specific attacks on our team's competence and members and tweets encouraging people not to support us. Another "black hat" tactic employed by Tag Team was the acquisition of a large number of fake Twitter followers, ostensibly purchased. Two days before the challenge began, their number of Twitter followers went from 37 to over 12,000, literally overnight. Presumably, their goal was to give an impression that they were the most promising team, and thereby harness the bandwagon effect. This, however, does not seem to have helped, as their following did not subsequently increase very much.

In summary, it appears that most competitors focused purely on social media, almost exclusively using Twitter to spread their message. This narrow strategy was not sufficient, not least because several teams came to be perceived as spammers. As a result, none of the other teams successfully located more than 1 out of the 5 suspects.

Our team was far more selective in its tweets and social media strategy, but also made sure to generate as much attention as possible in blogs and online news. Even a small comparative initial advantage in credibility and awareness can lead to eventual victory in this competitive domain. This phenomenon can be explained in the following way. A valid photograph of one of the targets has value, because there is prize money provided by the organizers. Hence, any participant in possession such a photograph needs to decide carefully which team to send it to. The first barrier for a team in receiving this information is for the participant to simply be aware of them. But secondly, and just as importantly, the participant must also trust the team leadership to give the promised rewards for such valuable information. Finally, the participant must have confidence in the likelihood of the team's victory –that is, an expectation that other participants would also submit their photos to that team.

Focusing on these qualities, and enhancing them through carefully considered incentives, allowed our team to get an early lead in finding suspects. We then made sure to publicize these early successes in real time on social media, which contributed to even greater levels of credibility and awareness. This positive feedback loop ensured that we were always ahead of the other teams.

Nevertheless, other teams ranked higher than us in Klout score[3]. This is a measure of influence and credibility across a user's social network obtained by data mining. This allows us to quantify the influence of each team's online presence (Table 1). Tag Team, the team that sent out numerous tweets and appeared to have purchased thousands of Twitter followers, ranked first by Klout score. This raises an interesting question: how important is that score for time-critical tasks and how good is it at filtering out attempts to engineer a higher score? It may very well be that the score works for more conventional online behavior, but not for a time-critical task such as the Tag Challenge. Developing a "time-critical Klout score" is therefore an interesting open question.

A final tactic employed against our team seems to have been direct attack. Between GMT 16:57 and GMT 17:00, four submissions were received, each for a different suspect, and all from the same user. The submitted files turned out to be copies of the same

[3] `klout.com`

Team's Twitter Handle	Klout score	Reach	Amplif	Style
@TagChallenge	32	126	6	socializer
@TagTeam_	26	54	4	conversationalist
(winner) @CrowdScannerHQ	21	38	3	explorer
@Tag_Challenge	16	20	3	dabbler
@TeamRave	10	1	2	observer

Table 1: Table reporting the Klout score, Reach and Amplif for the five teams participating to the challenge

executable file. Upon dissection, this executable contained references to various critical operating system libraries. This was assessed as, at best, an attempt to annoy us, or at worst, an attempt to make our systems vulnerable to an attack. The affiliation of the attacker is unknown, but malicious intent could not be ruled out.

5 Lessons Learned

The main goal of the Tag Challenge was to test the limits of social mobilization in a time-critical competitive environment. We were able to find only 3 out of 5 suspects within the allotted 12 hours, in contrast to the Network Challenge, where all of the balloons were found in under 9 hours. However, the Tag Challenge was arguably more difficult as the targets were not static and we needed to recruit participants dispersed across the globe with three separate languages. These considerations make the result of the competition — finding 3 out of 5 suspects — a testimony to the power of social media, which really pushed the boundaries of social mobilization.

Looking again at the fact that we found only 3 out of the 5 suspects, we may extract some conclusions about the difficulty of the challenge, compared to the Red Balloon challenge. Members of our team were geographically close to only one of the cities where suspects were present –Southampton is one hour away from London. Yet, we failed to find the London suspect. This could be a first lesson: geographical proximity does not guarantee anything, and that ours is really a de-localized technique. Being in 4 countries, across 2 continents, and having to deal with 3 different languages can also hinder the propagation of recruitment. But a comprehensive understanding of the role of these factors requires more systematic investigation than is afforded by a single run of the challenge.

Interestingly, social networks did not play an explicit role in our victory. While Facebook and Twitter are the most natural tools for social mobilization, we found it extremely difficult to promote our Facebook page and Twitter accounts (other teams seemed to face the same predicament). For instance, all of our efforts amounted only to a few dozen Facebook "likes" and Twitter followers. Nonetheless, these media were extremely important for at least two reasons. Firstly, they added to the credibility of our team by displaying the team's history in the form of wall posts and tweets. Secondly, the reach extended beyond the immediate fans and followers. While fewer than 50 people "liked" our Facebook page, the number of unique users (as measured by Facebook statistics) exposed to it before the competition exceeded 300, and we consequently received around 500 hits on our main team's Web site in the 24 hours leading up to and during the challenge. Furthermore, Facebook and Twitter enabled a broadcast-style communication to the most important people—the ones who explicitly expressed their interest. In particular, updates about our successes (in finding targets) were immediately posted to Facebook, Twitter, and our Web site, sending an encouraging signal that we were the strongest team.

The majority of the visitors who came to our Web site did so after reading a post about the challenge on Slashdot and comments we left in forums discussing the challenge. More

traditional forms of online media also played an important role in making people aware of the challenge and our team: we were mentioned in CNET, ZDNet, and by our respective universities' press teams.

It is also worth noting that 2 of the 3 people in our team who found a suspect had an existing interest or connection to crowdsourcing. The suspect in Bratislava was found by a computer science graduate interested in crowdsourcing. David Alan Grier, president-elect of the IEEE Computer Society, produces a regular podcast on crowdsourcing and was responsible for finding the suspect in Washington D.C. This feature of successful participants could simply come from the fact that most of the participants in our team had an existing interest in crowdsourcing, an affinity that is important to motivation in social mobilization [10]. Indeed, the majority of our post-competition survey respondents indicated interest in social mobilization and crowdsourcing.

Another interesting aspect is that, like in the financial markets where successful and profitable strategies are copied by investors over time and lose profitability, query incentive cascades [11], while novel in the balloon challenge, were copied in this challenge. An open question, again, is how this would play against the use of this strategy, if at all. Furthermore, if crowdsourcing is used extensively and indiscriminately, the large number of such search tasks of which a person will be aware could possibly result in information overload, to the detriment of any one task. In this case, it would be useful to study as to how tasks can be targeted at specific individuals, depending on their location, interests, mobility, network structure, and other relevant features.

The influence of competition on the efficiency of the search deserves some attention. Since we found every other suspect that others found, we can assume that the competition with the others did not harm the propagation of our network. However, a conclusive answer requires further investigation into the role of competition.

In Figure 5 we summarize the main lessons learned.

Main Lessons Learned:

1. *Social mobilization is both easy and hard:* We can, in principle, find people within 12 hours; diversity of country/language/culture plays against social mobilization.
2. *Social networks are not a silver bullet:* Gathering a large Twitter/Facebook following is hard, in which case, publicity through other online media (e.g., social news websites, online forums, press releases) becomes indispensable.
3. *Credibility matters:* It is important to build a profile as a credible team, which has good chances of winning.
4. *Incentives matter:* Focus on recruiting motivated people with existing domain knowledge, and affinity with the task.
5. *Effect of competitiveness is little understood:* Whether inter-team and/or inter-task competition hinders overall social mobilization warrants further investigation.

Figure 5: Summary of lessons learned

6 Concluding Remarks

Our success in the Tag Challenge is another testimony to the power of technology-mediated social mobilization. Involving vast geographical spread, target mobility, and time criticality, it has profound implications for social search. The challenge also helped us appreciate the limitations of the technology and the challenges it faces. The fact that we could not locate all targets within the allotted time highlights the difficulty of the task, and shows that social media is not a silver bullet. Much work needs to be done to exploit the technology's full potential, but also to put in place safeguards for preventing its abuse.

We foresee growing potential for applications that support time-critical social mobilization in the real world. In particular, one application is employing social networks for finding missing children or missing persons. Our experience complements recent lessons learned about the use of crowdsourcing in processing large number of satellite images in the search for computer scientist Jim Gray, who was lost at sea –a search that unfortunately did not succeed in locating him [7]. There is a need for more dedicated platforms for facilitating this, allowing efficient sharing of useful information. As is the case with the Tag Challenge and the Red Balloon Challenge, such applications will likely benefit from explicitly rewarding the routing of information and recruiting the right individuals for the search (e.g., those familiar with the geographies or urban areas of interest) [9]. One can also use these techniques to enable people to report on environmental violations (e.g., someone dumping pollutants in a river), or to map natural disasters in real-time (e.g., the spread of a forest fire) or after they occur (e.g., mapping damage after an earthquake) [8]. Again, social networks may be very useful here, since there is potentially useful information embedded within the network about volunteers that may have access to geographically relevant information. Crowdsourcing using social networks and human-based sensing can also be used to encourage data collection by citizen science volunteers [6]. For example, the Cornell Lab of Ornithology are already using volunteers to collect large-scale scientific data about urban and wildlife [3].

In short, this technology can be used to mobilize volunteers in faraway places to achieve highly distributed, possibly time-critical tasks. We are working on a general "social mobilization app" that anyone can use to build and coordinate teams of volunteers rapidly to address difficult, geographically distributed challenges [6].

Around 60 years ago, legendary social psychologist Stanley Milgram redefined our notion of social distance with his "six degrees of separation" experiment [14], showing that we are, on average, only 6 hops of friendship away from anyone else on earth. Facebook found the degree of separation to be only 4 in their network [1]. Endeavors like the "Tag Challenge" are set to redefine our conception of the temporal and spatial limits of technology-mediated social mobilization in the Internet age, showing that we can find any person (who isn't particularly hiding) in less than 12 hours.

Acknowledgments

We are grateful to the organizers of the Tag Challenge for running the experiment, and to all our volunteers, especially those who found the targets.

References

[1] L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna. Four degrees of separation. *arXiv:1111.4570*, 2011.

[2] J. Bagrow, D. Wang, and A. Barabási. Collective response of human populations to large-scale emergencies. *PLoS ONE*, 6(3):e17680, 2011.

[3] R. Bonney, C. Cooper, J. Dickinson, S. Kelling, T. Phillips, K. Rosenberg, and J. Shirk. Citizen science: a developing tool for expanding science knowledge and scientific literacy. *BioScience*, 59(11):977–984, 2009.

[4] M. Cebrian, L. Coviello, A. Vattani, and P. Voulgaris. Finding red balloons with split contracts: robustness to individuals' selfishness. In *Proceedings of the 44th symposium on Theory of Computing*, STOC '12, pages 775–788, New York, NY, USA, 2012. ACM.

[5] P. S. Dodds, R. Muhamad, and D. J. Watts. An Experimental Study of Search in Global Social Networks. *Science*, 301(5634):827–829, 2003.

[6] E. Hand. Citizen science: People power. *Nature*, 466(7307):685, 2010.

[7] J. Hellerstein and D. Tennenhouse. Searching for Jim Gray: a technical overview. *Communications of the ACM*, 54(7):77–87, 2011.

[8] M. Hesse. Crisis mapping brings online tool to Haitian disaster relief effort. *Washington Post*, January 16, 2010.

[9] D. Horowitz and S. D. Kamvar. Searching the village: models and methods for social search. *Commun. ACM*, 55(4):111–118, Apr. 2012.

[10] J. Iribarren and E. Moro. Affinity paths and information diffusion in social networks. *Social networks*, 2011.

[11] J. Kleinberg and P. Raghavan. Query incentive networks. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 132–141. IEEE, 2005.

[12] W. Mason and D. J. Watts. Financial incentives and the "performance of crowds". In *Proceedings of the ACM SIGKDD Workshop on Human Computation*, pages 77–85, New York, NY, USA, 2009. ACM.

[13] P. P. Meier. Verifying crowdsourced social media reports for live crisis mapping: An introduction to information forensics, November 2011. unpublished manuscript.

[14] S. Milgram. The small world problem. *Psychology Today*, 1(1):66–67, 1967.

[15] S. Milgram. The experience of living in cities. *Science*, 167:1461–1468, 1970.

[16] V. Naroditskiy, I. Rahwan, M. Cebrian, and N. Jennings. Verification in referral-based crowdsourcing. *PLoS ONE*, 2012. under revision.

[17] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland. Time-critical social mobilization. *Science*, 334(6055):509–512, 2011.

[18] J. Tang, M. Cebrian, N. Giacobe, H. Kim, T. Kim, and D. Wickert. Reflecting on the darpa red balloon challenge. *Communications of the ACM*, 54(4):78–85, 2011.

[19] D. Wilcock. Police 'facewatch' app targets london riot suspects. *The Independent*, 2012.

Biographies

Iyad Rahwan is Associate Professor and head of the Computing and Information Science program at Masdar Institute of Science & Technology, and Honorary Fellow at the University of Edinburgh. Contact him at `irahwan@masdar.ac.ae`.

Sohan Dsouza is a Research Engineer at Masdar Institute of Science & Technology, UAE. Contact him at sdsouza@masdar.ac.ae.

Alex Rutherford is a Post-doctoral Research Fellow at Masdar Institute of Science & Technology, UAE. Contact him at arutherford@masdar.ac.ae.

Victor Naroditskiy is a Post-doctoral Research Fellow at Southampton University. Contact him at vn@ecs.soton.ac.uk.

James McInerney is a PhD candidate at Southampton University. Contact him at jem1c10@soton.ac.uk.
UK

Matteo Venanzi is a PhD candidate at Southampton University. Contact him at mv1g10@ecs.soton.ac.uk.

Nicholas R. Jennings is Professor of Computer Science in the School of Electronics and Computer Science at Southampton University, and a Chief Scientific Advisor to the UK Government. He is a fellow of IEEE and the Royal Academy of Engineering. Contact him at nrj@ecs.soton.ac.uk.

Manuel Cebrian is a Senior Scientist at National ICT Australia, and Senior Research Fellow with the Media Laboratory at MIT. Contact him at cebrian@mit.edu.

Author Addresses

Dr. Iyad Rahwan
Masdar Institute of Science & Technology
PO Box 54224, Abu Dhabi
United Arab Emirates

Sohan Dsouza
Masdar Institute of Science & Technology
PO Box 54224, Abu Dhabi
United Arab Emirates

Alex Rutherford
Masdar Institute of Science & Technology
PO Box 54224, Abu Dhabi
United Arab Emirates

Victor Naroditskiy
Electronics & Computer Science
University of Southampton
Southampton SO17 1BJ, UK

James McInerney
Electronics & Computer Science

University of Southampton
Southampton SO17 1BJ, UK

Matteo Venanzi
Electronics & Computer Science
University of Southampton
Southampton SO17 1BJ, UK

Nicholas R. Jennings
Electronics & Computer Science University of Southampton
Southampton SO17 1BJ, UK

Manuel Cebrian
Optimization Research Group
Victoria Research Laboratory
Level 2, Building 193
Electrical and Electronic Engineering
The University of Melbourne, VIC 3010
Australia