Developing a Tool for Mega-Collaboration

Chris Newlon, IUPUI School of Informatics

Introduction

In 1997, Jakob Nielsen defined the term *mega-collaboration* as follows:

> "[T]he collective behavior of millions of people can form an environment where value is derived from the mass of actions even though each individual action is performed purely for the sake of the individual user." [1]

This term referred originally to a self-serving action such as linking to someone's Web site, thus unwittingly raising their search-engine rating. However, we are now seeing mega-collaboration in which the individual actions are deliberately aimed at resolving a common problem.

The need for a tool to facilitate large-scale collaboration has been illustrated by the public response to several recent disasters, such as Hurricane Katrina in New Orleans. What has emerged is a pattern where the formal disaster response of the government agencies is hobbled by rigid communication channels and conflicting organizational cultures while informal groups of private citizens coordinate across the Internet to provide much of the disaster relief.

Studies by the General Accounting Office, [2] and by Congress [3] after Hurricane Katrina illustrate this new phenomenon. What went wrong? The studies detailed a chaotic government response, with multiple chains of command and a myriad of approaches, where volunteers and donations were not well integrated. In examining why this happened, these studies described cultural barriers between agencies that impeded the flow of information. However, a review of news articles [4-6] reveals another story that may be even more significant – the story of what went right. Immediately after the hurricane, several private organizations and even one private individual, maintained on-line missing-persons databases to help people find their loved ones. Hundreds of thousands of Internet blogs, listbots, and bulletin boards provided highly specific information to people who needed it. [7, 8] Socially connected information networks channeled enormous amounts of money and resources into the recovery effort. [5, 9] In particular, organizations with local connections on both ends, such as neighborhoods, churches, unions, and professional associations, were highly effective.

The dynamic behind this effectiveness is something called the "small-world principle". Basically, it states that any two individuals in the network are likely to be connected through a short sequence of intermediate contacts. The effectiveness of small-world networks in moving information and resources to people in need is dependent on the amount of context they can provide. Research on the small-world principle by Kleinberg [10] demonstrated that a correlation between local structure and long-range connections is essential; otherwise individuals can't tell which contact is likely to lead to an efficient connection chain. This is why local connections at both ends yielded such effective targeting.

The implication of this is that the Internet has empowered individuals while centralized authority struggles to respond. Without some solution to the problem, this difference in functionality will only continue to grow. Therefore, some method is needed to interface the spontaneously forming mega-collaboration with the formal power structure. It was with this in mind that the specifications for this tool were developed.

While this project builds on research in various fields, it is most closely tied to work in disaster response inspired by Internet use during the large-scale disasters of the past two years. The emergence of "hastily formed networks" [11] or "ephemeral groups" [12] has characterized these disaster responses. Denning introduced the term "conversation space" to describe the medium used for communication during a disaster response. This conversation space consists of the physical systems by which people communicate; the players, including their roles, competencies, and authorities; and the "interaction practices" that the players follow to organize their cooperation. After examining the responses to both 9/11 and Hurricane Katrina, Denning observed, "One of our early conclusions was that the effectiveness of [hastily formed networks] rests on the quality of the conversation space established at the outset". [13] Its quality depends on participants agreeing on interaction rules and forming a consensus on the definition of the problem. This is the process of negotiation that the mega-collaboration tool is specifically designed to support.

A number of complementary interfaces for disaster response are beginning to emerge. Farnham and Kirkpatrick helped ephemeral responder groups implement Microsoft's Groove Virtual Office during the Katrina response. [12] This tool provided peer-to-peer networking and office function support for temporary operations facilities, such as field hospitals. Murphy created the Joint Victim/Responder Situational Awareness Decision Support System during the Katrina response. [13] This web-based application allowed victims to enter their needs, and responders to answer them. A group at Pennsylvania State University is working on "geocollaborative" crisis management, including a web portal to enhance group knowledge development through integration with external geospatial resources. [14, 15] Denning's group has implemented a prototype information exchange website [16] that is similar in concept to one proposed by Murphy during the Katrina response. [4] The proposed mega-collaboration tool would work in conjunction with these kinds of interfaces. It is specifically intended to help gather and organize data into formal structures that can be linked to more specialized software through open-standards interfaces.

In developing the concept for this tool, recent research in a number of different fields was drawn upon. One of these areas involves "serious games," or the use of a game-type interface to address real-world problems. Raybourn [17] recommended the use of such games in bridging between the organizational cultures of different groups, because they offer a safe arena in which to hammer out disagreements. Brandt and Messeter [18] suggested that framing collaborative activities in the format of a game improves idea generation and communication between stakeholders by shifting focus to the game, thus downplaying power relations and other factors that might inhibit brainstorming. Wikis and tweaking contests demonstrate the potential of other types of contribution interfaces by encouraging people to spontaneously donate their expertise and to correct each other's work. Major examples include Wikipedia [19] and MATLAB programming contests. [20]

Research in cultural negotiation and team dynamics has increased the understanding of cooperation, suggesting future tools for online collaboration. For instance, a number of studies have been conducted on virtual teams [21, 22] and "extreme teams" of up to several hundred members, typically seen in emergency response situations. [23] Artificial intelligence software is under development for extreme teams that respond to emergencies. [23, 24] Team interactions can also be modeled on a more basic level, however. Ess and Sudweeks [25] and Hewling [26] described how virtual group participants from different cultures engage in an ongoing process of negotiation to generate a new "third" culture through their unique interactions, where the "culture" is constructed out of the participants' online cross-cultural encounters. Barnum, [27] building on Tuckman's model of the team development process, [28] suggested that a team-building activity in the forming stage can help to jumpstart the team. This suggestion was corroborated in a controlled study documented by Farnham et al., [22] describing the use of a communication-support tool that superimposed a pre-written script on a simple Internet chat format.

In addition to all this, research on task-oriented data structures has increased the understanding of content representation, with implications for information management. In designing a data structure, almost any entity with which a team works can be characterized as an event, a goal, a task, a role, or an object, or resource. [29] A review of several task-based conceptual frameworks [29] reveals that at least two dimensions are necessary to describe these entities. In addition to the parent-child relationships between the various decomposed or grouped entities, on any given level the event, goal, task, role, and resource are all related to each other, in that the event triggers the goal, which "is had by" the task, which "is performed by" the role, which "uses" the resource. One problem that a spontaneously developed data structure must address, however, is that one person's event may be another person's goal. Therefore, it is quite important that any definitional structure be segregated from the physical phenomena that it describes.

Development Methods

The first step in designing this tool was to develop a set of general specifications, based on the research described above, and on a set of use case scenarios describing who would use the tool and why.

Cognitive walkthroughs were then conducted with a couple of groups who were assisting with the design using a hypothetical disaster scenario. These identified the problem definition task as an obstacle to the use of the tool by ordinary people. To overcome this, a problem-definition wizard was added that enables each team member to form an individual mental model of the problem at hand during the first 10 minutes of each trial. As the team members work, the tool uses their responses to add entries to a database, and then draws from the database to create individual treemaps representing the structures each team member has defined. A chat window is available at the bottom of the screen to allow the team members to communicate during any of the coordination stages.

The client side of this tool is written in JavaScript and AJAX, while the server side is written in PHP and MySQL. These open-source languages were chosen because they are freely available and will facilitate future examination and development of the tool within the open source environment. The tool is deliberately browser-based in order to make it accessible to the largest number of potential collaborators.

As soon as the first phase of development is completed, a usability study will be conducted on the tool, looking for issues that could interfere with research trials. Once the initial testing is done, and any problems found are corrected, the effectiveness of the tool will be examined in a research study. This second study will be conducted on-line. The results will be analyzed statistically to determine what effect use of the tool has on the group's ability to come to a consensus.

Results So Far

The database that supports this activity is very general. [30] A class diagram of part of the data structure is shown in Figure 1. While the data definition wizard encourages the team members to define their problem in terms of events, goals, tasks, roles, and resources, the database treats each of these definitions as an entity. The name and description of each entity is, therefore, added to the entity table. However, because one person's goal may be another person's task, event, role, or resource, the database also has a situation table. By linking most of the information about an entity to one or more situation-specific keys, it is possible to separate the identity of the entity from the situation in which it is being defined. This allows entities to be combined if they are found to be the same thing, without losing the situational differences between the two definitions. The database also has a relation table where the relationships between different entities can be defined. Finally, the database has list and item tables, which allow miscellaneous descriptive information to be entered about each entity. With the addition of a few tables to track teams, conversations, and the match process, the resulting simple generic database can be used to store very complex data.

A screen shot of the interface under development is shown in Figure 2. The screen is divided into an instruction area, a chat area, a wizard area that assists the players in describing their mental model of the problem at hand, and an area that allows them to compare their models with those of others, using hierarchical treemaps as a form of representation. The treemap form was chosen for its ability to represent data structure in a confined space. Clicking

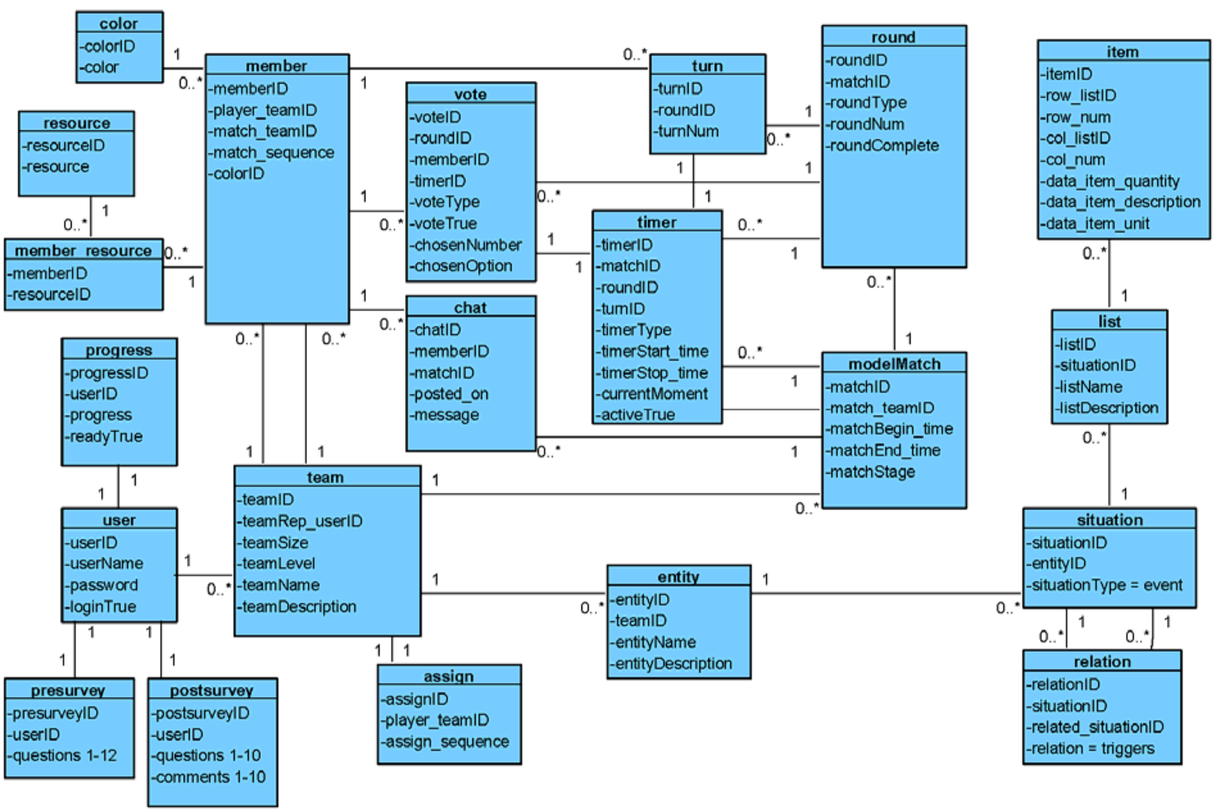

Figure 1. Class diagram of the data structure

on each cell of the treemap causes it to move to the next level of detail in the hierarchy, while clicking on the frame causes it to move back one level. Buttons below each treemap area allow the player to choose which player's treemap is displayed. Another set of buttons allows the player to choose which type of hierarchy to display, since the data can have three dimensions. These are the event—goal—task—role—resource sequence, the entity—list—item sequence, and the entity—sub-entity—sub-sub-entity sequence.

After initially spending ten minutes building individual mental models, the team members are next presented with a screen that displays each of the individual treemaps and a window where they can build their group treemap. The team members are given a few minutes to explore each other's definitions by clicking on various parts of these treemaps. The team members are then given a couple of minutes each, in succession, to work on the group data structure. They can use the wizard to add items to the treemap, or they can double-click on parts of the individual treemaps to add those parts to the group treemap. At the end of each round, the team is allowed to vote on whether they want another round, or whether they are ready to fill out an action schedule, based on their analysis of the problem. The successful completion of their action schedule is the final step in the team exercise.

Discussion

The project as described to this point has been developed as master's-level research. However, grants are currently being sought to continue development of this tool. The proposed project will build upon the work already completed by examining how the processes under development for within-team communication can be extended for between-team communication of sub-teams, allowing teams and topics of larger scale.

The issues to be tested include both team size and topic complexity. The number of possible interactions between pairs of people on a team is the factorial of the team size. Adding just one person increases the number of pair-wise interactions by the size of the team. When a team is small, it is beneficial to increase the number of these interactions because it adds to the richness of the conversation. But past a certain point, the team becomes unworkable. We must determine at what maximum size the team can continue to have a coherent online conversation and what size is optimal for disaster relief planning. We must also determine the best way to combine sub-team results if we divide the team and have the sub-teams work in parallel on the same topic.

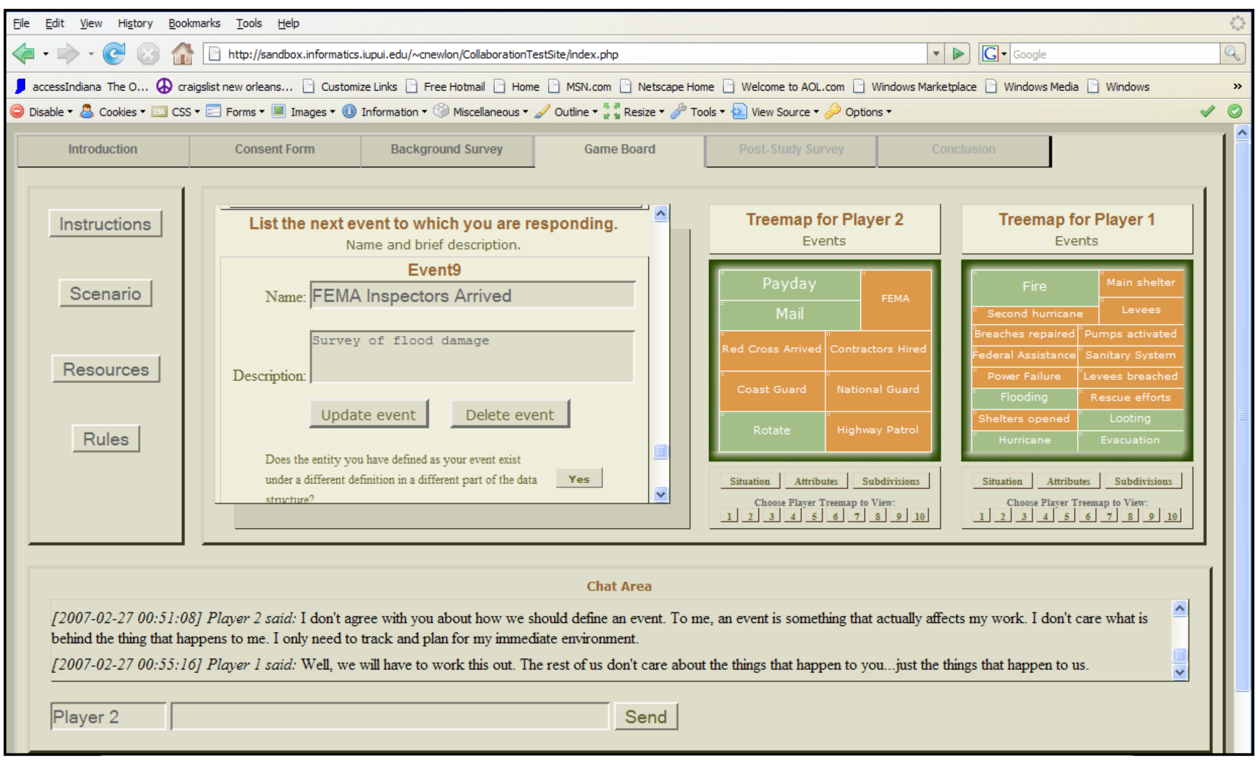

Figure 2. Screen shot showing the current version of the tool interface

While optimal team size will facilitate large-scale online collaboration, a more subtle problem that must be examined is topic complexity. An infinitely large team could theoretically address every topic that might be of interest to every team member. As soon as limits to team size are introduced, however, the topics a sub-team might address also become more limited. A solution is to have different sub-teams work on different topics. However, once both the team and the topic are divided, coordination between sub-teams suddenly becomes a more complex problem. Multiple sub-teams working on the same topic can share information and make decisions directly. Multiple sub-teams working on different topics must overcome greater definitional differences to coordinate their work. To solve this problem a playoff system is needed that is similar to the one used to determine team size to manage the complexity of the topic. It is here that the true power of the underlying data structure will be tested.

In theory, the value of this tool is that teams will avoid clashing and wasting resources but will instead find the most effective solution by pooling resources. However, there are questions about how this would work in real life. It may be very time-consuming. The representatives from dynamically-formed virtual teams may not have authority to dictate efficient use of resources, especially with respect to the break-up and recombination of physical teams. The combined plan may have some gaping holes that no one wants to address – no volunteers for the dirty jobs. Some conflicts may be hard to detect because of one plan "undoing" the pre-conditions of another plan. A high level of detail may be needed to avoid such problems. Most of these issues reflect problems that already arise from a multitude of uncoordinated grassroots efforts. The real question is whether, by allowing a structured interface between bottom-up spontaneous activities and top-down authority structures, this tool can solve these problems.

References

1. Nielsen J. Community is dead; long live mega-collaboration. Jacob Nielsen's Alertbox. 1997 [cited 2006, Feb 4]; available from: http://www.useit.com/alertbox/9708b.html.

2. Walker D. Statement by Comptroller General David M. Walker on GAO's preliminary observations regarding preparedness and response to Hurricanes Katrina and Rita. 2006 [cited 2006 Feb 4]; available from: http://www.gao.gov/new.items/d06365r.pdf.

3. Davis T, Rogers H, Shays C, Bonilla H, Buyer S, Myrick S et al. A failure of initiative - The final report of the select bipartisan committee to investigate the preparation for and response to Hurricane Katrina. 2006 [cited 2006 Apr 1]; available from: http://katrina.house.gov/full_katrina_report.htm.

4. A geek's-eye view of hurricane relief - autonomechs' Tim Murphy on what's being done, what needs doing. 2005 [cited 2005 Sep 8]; available from: http://www.msnbc.msn.com/.

5. Feulner E. Rebuilding in the Big Easy doesn't have to be so hard. 2006 [cited 2007 Mar 1]; available from: http://www.realclearpolitics.com/articles/2006/08/rebuilding_in_the_big_easy_doe.html.

6. Vijayan J. IT pros aid in search for Katrina victims - nonprofit groups, individual volunteers deploy technology to find the missing. 2005 [cited 2005 Sep 19]; available from: http://www.computerworld.com/.

7. Hurricane Katrina update 12-5-05. 2005 [cited 2006 Mar 24]; available from: http://www.bloggersblog.com/hurricanekatrina/.

8. Craigslist resources for Katrina survivors. 2005 [cited 2006 Mar 25]; available from: http://neworleans.craigslist.org/about/help/katrina_cl.html.

9. International Association of Fire Fighters. Reports from the hurricane frontlines: Katrina. 2005 [cited 2005 Feb 4]; available from: http://daily.iaff.org/Katrina/Katrina.htm?c=Report.

10. Kleinberg J. The small-world phenomenon: An algorithmic perspective. *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing.* 2000. p. 163-170.

11. Denning P. Hastily formed networks. *Communications of the ACM*, 2006;49(4):15-20.

12. Farnham S, Kirkpatrick R, Pedersen E. Observation of Katrina/Rita deployment: Addressing social and communication challenges of ephemeral groups. In *Proceedings of Intelligent Human Computer Systems for Crisis Response and Management*, Newark, New Jersey. 2006.

13. Murphy T. Joint Victim / Responder Situational Awareness DSS. 2005 [cited 2006 Dec 1]; available from: http://autonomechs.com/classes/IDS691/.

14. MacEachren A, Cai G, McNeese M, Sharma R, Fuhrmann S. GeoCollaborative crisis management: designing technologies to meet real-world needs. *ACM International Conference Proceeding Series.* 2006 [cited 2006 Dec 1]; 151:71-2. Available from: http://doi.acm.org/10.1145/1146598.1146624.

15. Tomaszewski B, MacEachren A, Pezanowski S, Liu X, Turton I. Supporting humanitarian relief logistics operations through online geocollaborative knowledge management. *ACM International Conference Proceeding Series.* 2006. p. 358-359. (Vol 151).

16. Denning P, Hayes-Roth R. Decision making in very large networks. *Communications of the ACM, 49*(11). 2006. p. 19-23.

17. Raybourn E. Computer game design: New directions for intercultural simulation game designers. *Developments in Business Simulation and Experiential Exercises,* 24. 1997 [cited 2005 Nov 15]; available from http://www.cs.unm.edu/~raybourn/games.html.

18. Brandt E, Messeter J. Facilitating collaboration through design games. *Proceedings of the Eighth Conference on Participatory Design: Artful Integration: Interweaving Media, Materials and Practices.* 2004. p. 121-131. (Vol 1).

19. Leuf B, Cunningham W. *The Wiki Way: Quick Collaboration on the Web,* Addison-Wesley. 2001.

20. Gulley N. In praise of tweaking. *Interactions.* 2004;11(3):18-23.

21. Powell A, Piccoli G, Ives B. Virtual teams: A review of current literature and directions for future research. *The Data Base for Advances in Information Systems*. 2004;35(1).

22. Farnham S, Chesley H, McGhee D, Kawal R, Landau J. Structured online interactions: Improving the decision-making of small discussion groups, *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*. 2000. p. 299-308.

23. Scerri P, Farinelli A, Okamoto S, Tambe M. Allocating tasks in extreme teams. *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems.* 2005. p. 727-734.

24. Scerri P, Xu Y, Liao E, Lai J, Sycara K. Scaling teamwork to very large teams. *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems.* 2004. p. 888-895. (Vol 2).

25. Ess C, Sudweeks F. Culture and computer-mediated communication: Toward new understandings. *Journal of Computer-Mediated Communication.* 2005 [cited 2006 Feb 15];11(1). Available from: http://jcmc.indiana.edu/vol11/issue1/ess.html.

26. Hewling A. Culture in the online class: Using message analysis to look beyond nationality-based frames of reference. *Journal of Computer-Mediated Communication*. 2005 [cited 2006 Feb 15]; 11(1). Available from: http://jcmc.indiana.edu/vol11/issue1/hewling.html.

27. Barnum C. Building a team for user-centered design. *Proceedings of IEEE Professional Communication Society International Professional Communication Conference and Proceedings of the 18th Annual ACM International Conference on Computer Documentation,* 2000. p. 325-332.

28. Tuckman BW. Developmental sequence in small groups, *Psychology Bulletin*. 1965;63:384-399.

29. van der Veer G, van Welie M. Task based groupware design: putting theory into practice. *Proceedings of the Conference on Designing Interactive Systems: processes, practices, methods, and techniques.* 2000. p. 326-337.

30. Newlon C, Faiola A. Support for mega-team collaboration with cultural tools: Providing a framework for the dynamic development of team culture. In F. Sudweeds, H. Hrachovec & C. Ess (Eds.), *Proceedings of Cultural Attitudes towards Communication and Technology*. Tartu, Estonia. 2006. p. 235-254.