A New Model for Mega-Collaboration

Christine Newlon
IU School of Informatics
535 West Michigan St., IT 475
Indianapolis, IN 46202 USA
cnewlon@iupui.edu

Karl F. MacDorman
IU School of Informatics
535 West Michigan St., IT 487
Indianapolis, IN 46202 USA

Paul Scerri
Robotics Institute, CMU
5000 Forbes Ave.
Pittsburgh, PA 15217 USA
pscerri@cs.cmu.edu

ABSTRACT
Public use of information and computer technology (ICT) in disaster response has become a new type of "mega-collaboration." However, government policy is moving in the opposite direction, toward a strict chain-of-command model. The resulting divergence can lead to inefficiency and weakened disaster response. Hence, there is a clear need for an interface that can bridge the gap between these two approaches. This paper presents a prototype solution as a focal point for discussing the application of current information and technology theory to the design of an interface for solving this problem.

Author Keywords
Mega-collaboration, disaster recovery, human-computer interaction, computer-mediated communication, mental models

ACM Classification Keywords
H5.3 [Information interfaces and presentation]: Group and Organizational Interfaces—collaborative computing, computer-supported cooperative work, theory and models, web-based interaction; I2.11 [Artificial intelligence]: Distributed artificial intelligence—intelligent agents; K4.1 [Computers and society]: Public policy issues—human safety.

INTRODUCTION TO MEGA-COLLABORATION
Nielsen first used the term "mega-collaboration" in 1997 to describe a collective behavior that is individually motivated but results in an outcome that benefits society [7]. The term originally referred to actions like unwittingly raising a website's search-engine rating by linking to it. However, recent disasters have spurred a new kind of mega-collaboration, where people from around the world decide to work together to respond to a crisis [6]. Palin and Liu [8] studied this new phenomenon and noted that information

and computer technology (ICT) has expanded the role of the public in disaster response. They described mega-collaboration as "an emerging form of societal-scale supported cooperative activity that extends and challenges our knowledge of computer-mediated interaction" (p. 727). Using the global reach of the Internet, large numbers of people, with a wealth of skills, time, and resources can come together to respond to events with unprecedented speed and commitment.

If the energy, skills, and resources of governments, NGOs and individuals could be effectively harnessed, the impact of disasters could be dramatically reduced. Unfortunately, with no way to dovetail official and spontaneous activity, this new grassroots empowerment could add to the general chaos of a disaster instead of reducing it [6]. Rather than addressing this problem, US government policy for formal disaster response appears to be running in the opposite direction. The US government has mandated the use of a quasi-military organizational structure for disaster response [8]. This protocol, the National Incident Management System (NIMS), routes all response through a single, unified command structure, which has the advantage of making each unit's responsibility clear. However, an analysis of the aftermath of Hurricane Katrina illustrates the hazard of organizing a response effort in this way; it is vulnerable to failure at a single point [2][14]. Even assuming no overt failure, this model works poorly in situations with many victims or volunteers [8]. Now that ICT has empowered the public, the conflict between these two approaches has become more conspicuous [6][8]. Palin and Liu call for designs to enhance the effect of citizen-generated information on the work practices of formal response organizations, to extend HCI/CSCW research to the improvement of command-and-control functionality in disaster situations [8].

A major impediment to effective coordination is that different participants work in dramatically different ways. For example, government and military responses use top-down planning and control, while grassroots volunteers organize and react in a bottom-up manner. This fundamental difference leads to inefficiency, interference, or even deadlock when these different groups try to work together.

We propose a mega-collaboration model as an over-arching framework helping these groups collaborate. The central

thesis of this model is that the human response effort can be divided into dynamically-populated sub-teams with the aid of web-based software agents. Each sub-team can develop its own model to define its part of the problem. Sub-team representatives can then consolidate these models in agent-facilitated compare-merge "playoffs," thus enabling large teams to agree on the nature and details of the problem and coordinate effective action. The information developed by this method can also be dynamically organized into a knowledge base, linking the collaborative activities of the public response to the command-and-control activities of the formal response.

We developed a prototype interface based on this conceptual model to provide a focal point for discussing the application of current ICT theory to novel interface design [5]. The prototype lets a single team to develop a simple group model and action plan while communicating via a chat interface. It has enabled us to explore future usability issues.

BACKGROUND OF THE DESIGN PROJECT

The mega-collaboration design project builds on research in several fields, particularly in agent-facilitated emergency response [10][11][12], and has been informed by research on Internet use during recent disasters. The emergence of "hastily formed networks" [3] or "ephemeral groups" [4] has characterized these disaster responses. Denning introduced the term "conversation space" to describe the medium used for communication during a disaster response. After examining the responses to both September 11th and Hurricane Katrina [3], 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" (p. 17). This quality depends, in part, on participants agreeing on interaction rules and forming a consensus on the definition of the problem. The mega-collaboration tool must be designed to support this process of negotiation.

THE PROPOSED INTERFACE

Our particular approach to mega-collaboration support is to develop a standardized negotiation protocol that a software agent can use. Our working hypothesis is that by enabling software agents to facilitate the negotiation process, human coordination can be scaled up exponentially. Furthermore, by dividing the effort among sub-teams of people and creating an autonomous agent to coordinate each sub-team's interaction with other sub-teams (Fig. 1), we hope to create a multi-dimensional matrix wherein a symbiotic mixture of human and agent initiative leads to large-scale collaboration.

Capturing, storing, and providing visualizations of teammates' mental models are essential to large-scale collaboration. To determine the potential of such a tool, we created a prototype interface in AJAX, PHP, and MySQL [5] for use as a test bed. This prototype interface (Fig. 2) currently supports the negotiation protocol and mental model formation with a chat window, an expanding form for entering structured data, and a treemap for visualizing the data. The agent program determines the optimum number of teams to create from each group, based on the number of users logged in, and manages the flow of the exercise. Each teammate is encouraged to develop an individual model of the problem, to compare the model with those of other teammates, and to negotiate the consolidation of all their models into a team model. The team then develops an action plan based on this consolidated model.
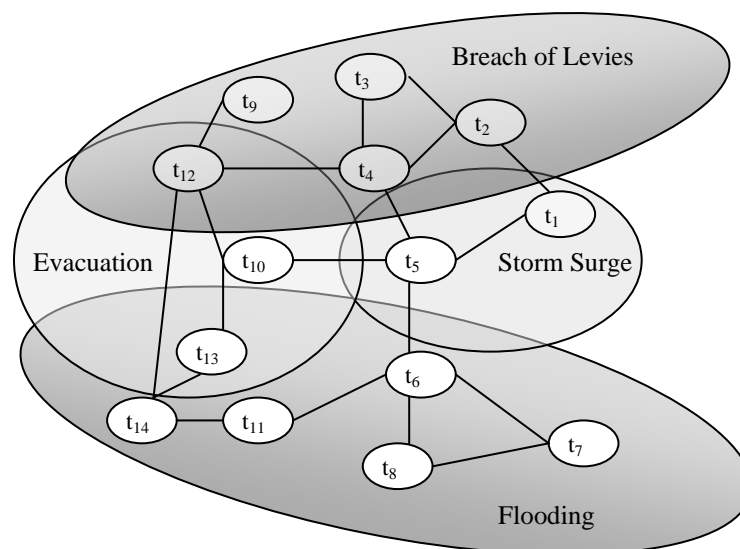

Figure 1. The relationship between topic branches and agent-managed coordination of sub-teams shows teamwork among agents, each of which is representing teamwork among humans. In effect, each human sub-team becomes a method of an agent object.

Placing intelligent agents at the core of the system will open up a range of possibilities for significantly improving mega-collaboration. For example, agents can constantly monitor details, proactively identifying potential conflicts or synergies, bringing them to the attention of relevant sub-teams. Agents can also monitor information needs and route newly-produced information rapidly. Proactive agents and infrastructures can automatically track activities and automate routine coordination, dedicating more human effort to the response and less to its administration.

Carnegie Mellon University's small-world networking architecture has been successfully used to test the goal coordination of large agent teams facing an emergency response scenario, and to allocate roles and tasks to these teams [10][11][12]. Therefore, it is worthwhile to explore a similar approach to combine agent and human teamwork in increasing the capacity of agent-managed matrices.

In extending this research, agents will be created to represent their own sub-teams of humans. These agents will establish a network of associations with each other based on the similarities being demonstrated by their human sub-teams. When a conflict is detected, the respective agents will arrange a compare-merge playoff to resynchronize the relevant portions of each sub-team's world model and resolve the conflict. Suppose two sub-teams plan to evacuate the same church. When they send representatives to the playoff team, each group will bring a data structure identifying the church, the goal of evacuating it, and other related information. Because both user sub-teams and

official entities will have detailed their approaches to the problem as part of their modeling process, the agents can integrate the solutions of the different groups to generate a suggested amalgamation automatically. The planners who were once at odds will be assisted in reuniting and realigning by their respective agents.

As part of model consolidation, the details from each model will be combined, and any duplicate items will be eliminated. The playoff teammates can then negotiate via the chat room. Notice that while many applicable agent technologies are already available, their use has been limited because of an inability to interface with real human organizations. Our prototype aims to address this technology gap.

BEYOND THE PROTOTYPE
Several problems must be solved to move from the small test bed currently in use to a useful tool for a real emergency.

Teaming
The first challenge facing mega-collaboration in the context of disaster response is the team formation procedure. One model already in use is the spontaneous teaming of massively-multiplayer online role-playing games (MMORPGs) [9]. Typically, someone will pick a quest or other goal and then issue a call to the general chat room. Others who want to participate will answer. The nascent team then establishes its own private conversation. Another
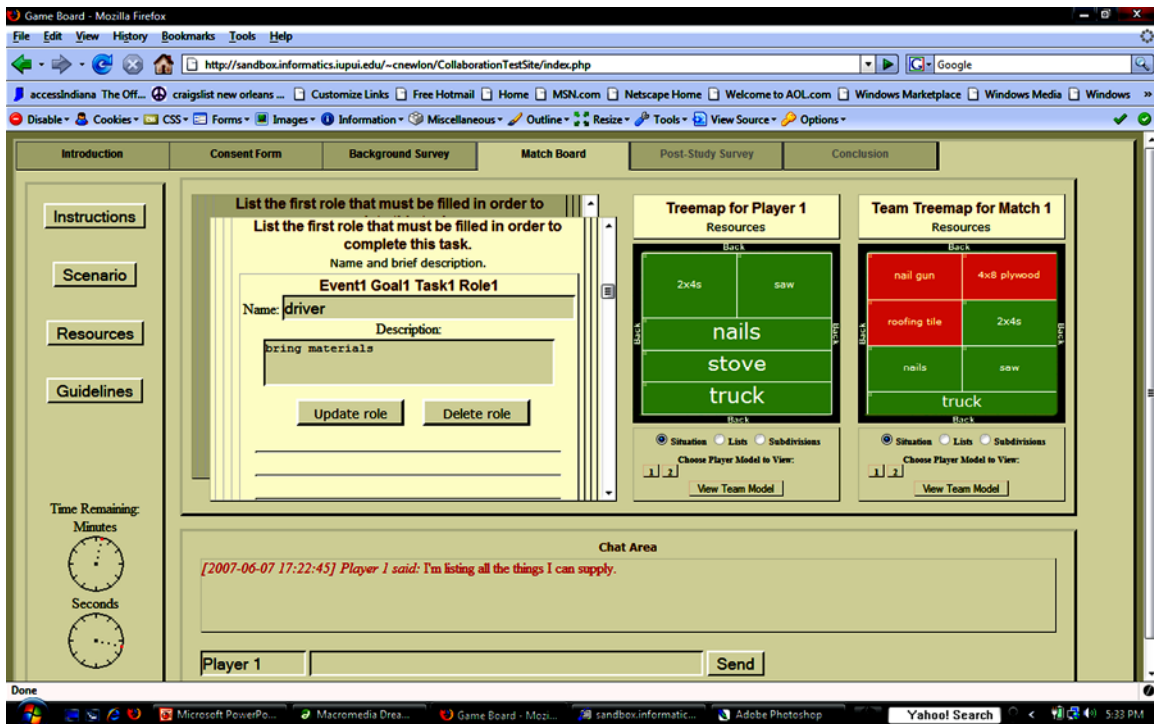

Figure 2. This screen shot shows the prototype interface for virtual collaboration. It features an expanding data entry form, treemaps representing portions of the resulting database, a chat window, and various navigation buttons.

possibility is to have the tool itself form teams by random assignment, the use of friends' lists, or proximity. Alternatively, participants could scan the topics and teams already established and join a team according to their preference. The final interface should probably be a synthesis of all of these. In addition to these unrestricted teaming strategies, allowances must also be made for the formation of restricted participation teams, with assigned participants, as would be required by a formal chain-of-command structure.

Input Interfaces

Usability testing of the current test bed has helped to determine the interfaces needed for model development. It was evident from the beginning that several formats would be needed for input. As of this writing, the clear winners are data-entry tree, data grid, and list entry, all of which have standard widgets available in various development environments. Another envisioned interface would enable participants to "felt-board" their hierarchies by moving disconnected tree structures around as the situation's structure becomes clearer. This is similar to the white-board pattern.

In addition to these general data entry formats, enabling participants to customize formats would help structure their thinking while boosting their creativity. In particular, action plan items should be part of the model, rather than a separate interface, as they are in the current test bed. Other input forms for scheduling, contact information, and location information would also be useful. It would be a leap in functionality to allow participants to develop their own specific input formats that could be adapted to the problem at hand and made available to others.

Output Interfaces

Despite the obvious need for multiple display interfaces, the current test bed has only a treemap. Although the treemap was not well received in its prototype form, it is nevertheless the most space-efficient form of display available. Our plan is to enhance its visualization, displaying each cell's contents instead of just listing the number of items in each cell.

Participants specifically requested additional interfaces, such as a flowchart of the model. A flowchart (Fig. 3) has been the only effective way to illustrate the models that resulted from runs of the current test bed. Currently, these are created by hand from the database. Hence, it is important to generate flowcharts automatically. A disadvantage of flowcharts is that they are space inefficient, so scrolling is often required. The data structure generated by the tool will be a network, not a simple tree. The assumption is that viewing "chunks" of the network as hierarchies can improve both development and display. People tend to think in this way. A method for navigating between chunks would need to be determined.

In addition to the general display formats, specific display formats for specific types of data will be needed. These will probably mirror the specific entry formats, namely, action plans, schedules, address books, maps, and so on. The greatest leap in functionality will come from the participants' ability to sort and search the available data to create the specific output in the specific format they want and then to make this information available to whomever they choose. This is true whether the participant is a member of a spontaneous team or a formal chain-of-command team.

Model-Building Process

The current test bed uses a rigidly-timed process that encourages the formation first of individual models, then of a team model, then of an action plan. Experience from testing indicates that beginners need a period of training on the interface, but that formation of individual and team models tends to happen concurrently, along with chat room discussion of the problem. Some standardization will be necessary to interconnect the different teams, but it should be possible for each team to set up its own timers and to call its own votes.

Interconnection (Playoff) Agents

The mega-collaboration concept hinges on the scalability of teams. Someone should be able to visually inspect the tool's output and determine that several teams need to get together and compare their models. However, with all users concentrating on their own piece of the problem, this comparison may be difficult, so the process should be assisted by the agents. The power of having all the activity take place in one data structure is that each forming team can automatically generate its own agent that searches for other teams developing similar models and coordinates compare-merge playoffs with the agents of those teams. Given the network structure of the data, these playoffs could take place in any direction, though an operational hierarchy will probably develop that mirrors the hierarchy of the relief effort.

Although the playoff concept has always been integral to the conceptual framework of the prototype, the current test bed focuses on the interfaces that connect a single team and that team's ability to choose a team representative. The agents, the search procedure, and the playoff process all remain for the next phase of development.

Data Structure: The Big Picture

Although the current database structure (Fig. 4) will need to be elaborated as features are added, it has tested very well and appears to be adequate in concept to support the tool. The current test bed makes use of a prewritten scenario for an initial crisis. However, a working version of the tool will need to draw from the developing data structure to provide scenario information to the participants. One way to do this would be to let the participants browse the data structure on their own to form a mental picture of the problem. Another

Events

Robinson
patch roof damage

Jones house
repair kitchen and siding

Smith house
repair damaged roof

Goals

Transport
move materials to the site

Gather materials for job
o 1000 board-ft of lumber o 20 sheets of plywood o 9 squares of roofing o 1000 sq ft of siding o 500 ft of electrical wire o 50 ft of plumbing pipe o 40 person-hours of labor o 1 refrigerator o 1 stove

Gather materials for job
o 100 board-ft of lumber o 50 sheets of plywood o 16 squares of roofing o 50 ft of electrical wire o 40 person-hours of labor o 3 beds o children's clothing

Deliver Materials
get the stuff there

Deliver Materials
truck it there

Start Construction Work
get the people together & assign a leader

Tasks

Gather supplies
* 20 board-ft of lumber * 3 sheets of plywood * 2 squares of roofing * 15 person-hours of labor * man's size medium clothing

Move materials to site
drive truck to site

Roles

Driver
someone at the wheel

Figure 3. Diagram of team model developed using the Mega-Collaboration Prototype Tool.

way would be to develop a "scenario-building" process, which could be agent-driven, that abstracts from the data structure and provides drill-down links.

The developing structure will be at varying levels of maturity, depending on how many playoffs each of its parts completed. This maturity information must be maintained for each data item, because it is a measure of confidence concerning the accuracy of the data. Counting the number of edits and links to other models could be sufficient.

Another feature of the data structure is access control. The test bed already controls who can edit and view items at any given time. If this tool is to interface with chain-of-command organizations, it must also enable restricted participation teams to control access to their models. This control should extend to the data-item level, so that such teams can selectively release information. One issue is how much access the agents will have. Although the restricted participation team's agent will be able to identify potential conflicts in the models of unrestricted teams, what about

conflicts with restricted teams? Security clearance levels may be required, even for the agents. In a global situation, with multinational restricted teams interacting with unrestricted teams, some warning system may need to be developed. Another challenge is conducting a playoff when one or more of the teams is restricted. Presumably, each restricted team could alter its model based on any unrestricted information available, but the playoff model could only use information a restricted team had specifically released.

Interface with the Chain-of-Command
This tool provides a place where individuals and groups make explicit, step-by-step plans and store information generated about the unfolding disaster and response. Intelligent software agents use this formalized information environment to integrate and streamline activities not only within a particular coordination paradigm but, ultimately, across paradigms.

Information on volunteer activities and resources available constitutes the greatest benefit of this tool to chain-of-command agencies. As mentioned, the ability to construct output reports and determine the maturity of each data item would be an important part of the tool's functionality for all users, not just formal organizations. However, the tool would provide an interface for negotiation between the two

different organizational cultures in addition a knowledge-base that both could use.

We propose matching the teams according to similarities in the objects they are modeling and having the agents manage the negotiation process. However, the agents can also track the authority structures of the teams based on how they are formed. Therefore, in theory top-down formulated rules can be added to the negotiation process for restricted-access teams.

Another possible interface for chain-of-command organizations would be an output report resembling a request for proposal. Once the government formed a top-down plan, the generation of a formal list of requests would allow agents to connect their teams to relevant problems selectively. This process could be followed by any organization seeking to coordinate through a mega-collaboration framework.

Notice that the interface supports both bottom-up and top-down information flows. Bottom-up organizations can be alerted to opportunities or holes in a top-down plan, showing where their efforts might be most effective. The tool can give them rapid access to all the information available in the entire response and ensure that local efforts work synergistically with the larger response. This access should encourage volunteers to use the tool.
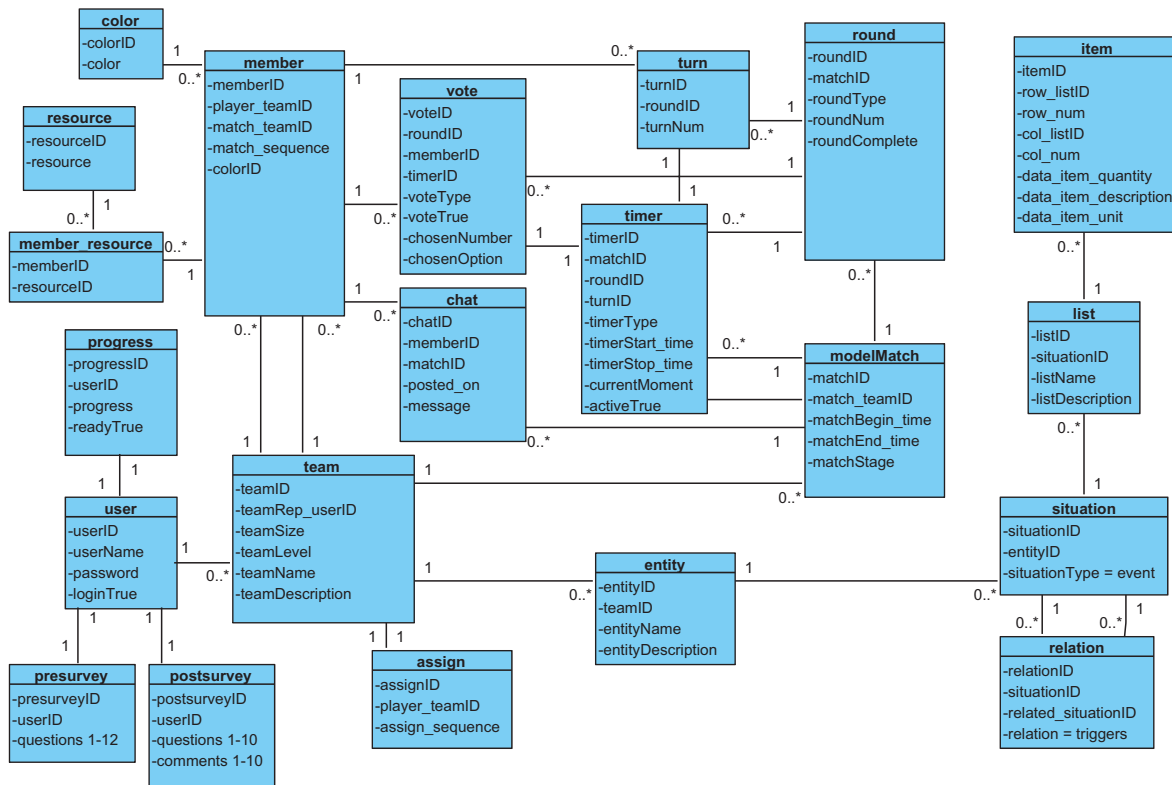

Figure 4. Class Diagram of Data Structure

Implications of Theory

This project has reached a stage at which design ideas can be drawn from current ICT theory and tested using the prototype. For instance, Crapo, in his cognitive-theoretic survey of visualization and modeling [1], recognizes the problem identification stage of model development as separate from the problem definition and structuring stage [1]. The prototype design instructs the teammates to begin by working individually on problem definitions. This has been difficult for the users, however. Not only do they want a separate problem identification stage before the definition stage, they want it to be a group activity conducted in chat. Apparently, teammates need to become familiar with their group context in a social setting even before developing their individual definitions of the problem.

Another theory that has major implications for the success of such an interface is terror management theory. It has been demonstrated that individuals who have been reminded of their own mortality (as will be the case in most disaster situations) tend to cling to their cultural worldviews, look for strong leaders, and display more hostility toward out-groups or perceived external threats [13]. However, a structure that enfolds many different organizational entities can potentially convert them from external threats to internal resources in the minds of those involved in disaster response. By providing a flexible, but structured interface for negotiation and dialogue, such a tool could facilitate development of the culture and leadership needed to respond to the crisis.

CONCLUSION

Mega-collaboration is an emerging phenomenon as the public begins to use ICT in responding to disasters. The US government is moving in the opposite direction, however, emphasizing a strict chain-of-command model. There is a need for effective interfaces to support mega-collaboration and connect it to chain-of-command structures. This paper introduced our prototype as a focal point for discussion on applying current ICT theory to the conceptual development of such an interface. The ideas we contribute to this discussion include iterative scaling, negotiated mental models, and the use of autonomous agents.

REFERENCES

1. Crapo, A., Waisel, L., Wallace, W., & Willemain, T. (2000). Visualization and the process of modeling: A cognitive-theoretic view. In *Proc. KDD* (pp. 218-226).

2. Davis, T., Rogers, H., Shays, C., Bonilla, H., Buyer, S., Myrick, S. et al. (2006). A failure of initiative: The final report of the select bipartisan committee to investigate the preparation for and response to Hurricane Katrina. http://katrina.house.gov/full_katrina_report.htm.

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

4. Farnham, S., Kirkpatrick, R., & Pedersen, E. (2006). Observation of Katrina/Rita deployment: Addressing social and communication challenges of ephemeral groups. In *Proc. Intelligent Human Computer Systems for Crisis Response and Management.*

5. Newlon, C. Developing a Tool for Mega-Collaboration. (2007). http://archive.nmc.org/events/2007summerconf/materials/Mega-Collaboration.pdf.

6. Newlon, C., & Faiola, A. (2006). Support for mega-team collaboration with cultural tools: Providing a framework for the dynamic development of team culture. In *Proc. Cultural Attitudes towards Communication and Technology* (pp. 235-254).

7. Nielsen, J. (1997). Community is dead; long live mega-collaboration. Jacob Nielsen's Alertbox. http://www.useit.com/alertbox/9708b.html.

8. Palen, L. & Liu, S. (2007). Citizen communications in crisis: Anticipating a future of ICT-supported public participation. In *Proc. CHI* (pp. 727-736).

9. Rauterberg, M. (2003) Determinants for collaboration in networked multi-user games. Ryohei Nakatsu & Junichi Hoshino (Eds.), *Entertainment Computing: Technologies and Applications*. IFIP Kluwer Academic Publishers, 313-321.

10. Scerri, P., Farinelli, A., Okamoto, S. & Tambe, M. (2004). Allocating roles in extreme teams. In *Proc. Third International Joint Conference on Autonomous Agents and Multiagent Systems, 3*, 1502–1503.

11. Scerri, P., Farinelli, A., Okamoto, S., & Tambe, M. (2005). Allocating tasks in extreme teams. In *Proc. Fourth International Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 727–734).

12. Scerri, P., Xu, Y., Liao, E., Lai, J., & Sycara, K. (2004). Scaling teamwork to very large teams. In Proc Third International Joint Conference on Autonomous Agents and Multiagent Systems 2 888-895.

13. Solomon, S., Greenberg, J., & Pyszczynski, T. (1998). Tales from the crypt: The role of death in life. *Zygon: Journal of Religion and Science, 33,* 9-43.

14. Walker D. (2006). Statement by Comptroller General David M. Walker on GAO's preliminary observations regarding preparedness and response to Hurricanes Katrina and Rita. http://www.gao.gov/new.items/d06365r.pdf.