Community-Based Cloud for Emergency Management

Juan Li[1], **Qingrui Li**[1], **Samee Ullah Khan**[2], **Nasir Ghani**[3]
[1]Computer Science Department, North Dakota State University
[2]Electrical and Computer Engineering Department, North Dakota State University
[3]Electrical and Computer Engineering Department, University of New Mexico
j.li@ndsu.edu, qingrui.li@ndsu.edu, samee.khan@ndsu.edu, nghani@ece.unm.edu

Abstract - *Natural and man-made emergencies pose an ever-present threat to the society. In response to the growing number of recent disasters, such as the Indonesian volcanic eruption, Gulf of Mexico oil spill, Haitian earthquake, Pakistani floods, and in particular, the Red River crest that causes flood almost every year here in Fargo, North Dakota, we propose a community-based scalable cloud computing infrastructure for large-scale emergency management. This infrastructure will coordinate various organizations and integrate massive amounts of heterogeneous data sources to effectively deploy personnel and logistics to aid in search and rescue. The infrastructure also will aid in damage assessment, enumeration, and coordination to support sustainable livelihood by protecting lives, properties and the environment.*

Keywords: Emergency management, cloud computing, community computing, information integration, social networking.

1 Introduction

Natural and man-made disasters require an effective and efficient management of massive amounts of data and coordination of wide varieties of people and organizations. The data involved in emergency management includes geographical data about the infected area, data about shelters and available transportation means, data about victims and relief personnel, available rescuing resources, and measurements from the field. The data may belong to multiple autonomous organizations, such as government organizations (GO), non-governmental organizations (NGOs), international non-governmental organizations (INGOs), individuals, communities, and industries. Therefore, besides integrating and manage data from these different organizations, we also need to coordinate these organizations by enabling efficient communication and collaboration. Moreover, with the development of Web 2.0 technologies, such as Social Network Sites (SNS), blogs, wikis, and video-sharing, the general public is able to interact or collaborate with each other in a social media dialogue as creators of user-generated content in a virtual community. For example, there are growing trends of using social media to create entities in popular SNSs, such as Facebook, Twitter, and LinkedIn. At the same time, many emergency/disaster-related organizations, such as Federal Emergency Management Agency (FEMA) and Maine Emergency Management Agency (MEMA) have been set up at Facebook. Meanwhile, many hospitals have Twitter accounts for residents to follow their status in case there is an emergency. Expert groups have also been established on LinkedIn to facilitate emergency planning and wellness exercising techniques. Furthermore, the ubiquity of mobile wireless devices facilitates the general public's involvement in the generation, propagation, and consumption of information – the anywhere and anytime paradigm.

To manage emergencies involving large amount of data and wide varieties of organizations, we propose a cloud computing platform to provide on-demand, scalable and reliable compute services. According to The National Institute of Standards and Technology (NIST) [2], "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." Cloud computing can provide data/communication-as-a-service solution to emergency management.

Using classical cloud computing, tasks (i.e., data management and communication) are moved off the servers of agencies, such as fire, police, medical, etc, and to centralized data centers that can be accessed through the Internet. However a single data center is vulnerable to natural and man-made sabotage, especially in a disaster/emergency situation. For example, "the Federal Emergency Management Agency (FEMA) has been without access to years' worth of lessons-learned data for nine months, unable to recover access to it since a server failure in May 2010" [3]. In addition, moving large amount of data to a central data center is expensive. Moreover, data updating and synchronizing is difficult to achieve between the owner organizations and the data centers. Furthermore, transmitting data to a third-party location may cause many security issues.

The above factors motivate us to propose a distributed community-based virtual data center without the need for the involvement of a single data center. Our community-based cloud-computing solution will allow the fire, police, state government, emergency centers and other public agencies to connect their data sources, extending their access not only to others within their agency, but also to neighboring agencies, when authorized. According to Hilley [1], "The most important trends underlying all cloud offerings are 1) ubiquitous network connectivity and 2) virtualization." Hence our community-based distributed cloud commuting platform also follows these two principles. The decentralized architecture makes connectivity more robust by providing multiple accesses and redundancy to the cloud. A Service-Oriented Architecture (SOA) is used to virtualize resources from different contributors. It applies abstraction to unified data access and communication, decoupling resources' interfaces from underlying low-level details. The abstraction afforded by virtualized services means resource users do not need to be aware of sharing or cooperate explicitly.

The rest of the paper is organized as follows. Section 2 presents the related work to our propose platform. Section 3 details our system design. In Section 4, we evaluate the proposed method with a comprehensive set of experiments. Concluding remarks are also provided in Section 5.

2 Related work

During the recent years, information and communication technologies (ICT) have been applied for facilitating the management of natural disasters and man-made crises [5-7]. ICT can be applied during different stages of a disaster, including disaster prevention, mitigation, preparedness, disaster response, and disaster recovery. Basically, IT technologies can be used for: (a) effective warning of disasters using different communication channels, (b) integrating information on necessary supplies and other sources; (c) coordinating disaster relief work; (d) encouraging social, institutional, and public responses; and (e) evaluating the damages caused by a disaster and the need for disaster relief.

The development of ICT over the last few years has facilitated disaster management with numerous collaborative tools at different levels. In particular, some open source disaster management tools have become very popular [8], e.g., Ushahidi [9], Sahana [10], and SwiftRiver[11]. Ushahidi was developed to report on the violence during the 2008 Kenyan general election. The idea behind the Ushahidi system is to utilize the crowd sourcing information to enable the sharing of varieties of information in situations that cause numerous rumors and uncertainties. Since then, Ushahidi has been deployed more than 20 times around the world to cater for similar situations where little or no support is provided by governmental authorities responsible for emergency management. Sahana is "a web-based collaboration tool that addresses common coordination problems during a disaster, such as finding missing people, managing aid, managing volunteers, tracking disaster relief camps and the victims"[10].

As more and more people use social networking tools to communicate and collaborate, these tools are becoming even more important to organizations that provide emergency management services. Many works have reported how social networking tools can assist disaster management [12-14]. Based on the aforementioned studies, disaster and emergency management agencies are finding that social networking technologies may facilitate personal networking and massive citizen outreach. Moreover, social networking technologies provide good venues for getting feedback from the public (via Facebook and Ning, for example), locating subject-matter experts (via LinkedIn), and for communicating with communities large and small (via Twitter and wikis). Social media and social networking tools are, and will continue to be, important tools in the emergency manager's toolkit. In summary, the aforementioned studies have reported that ICT are critical for disaster management. Classical ICT and social media can assist in disaster management, early warning, disaster relief, and integrating societal assistance during a crisis (or a disaster event).

The more recently appeared Integrated Emergency Management (IEM), such as the European Access to Knowledge through the Grid in a mobile World (Akogrimo) project [4], sets up an environment in which the contributions of multiple agencies can be planned and coordinated in a coherent way.

3 System design

3.1 System overview

The proposed project aims to utilize the massive amount of resources inside and outside a community including humans, organizations, computing devices, and information to construct a collaborative social cloud computing architecture to support community disaster management. The architecture (see Fig. 1) of our proposed system consists of a web-based social network server that provides a platform to enable users (workers, first responders, local disaster related non-profit organizations, volunteers, and local residents) to access information, communicate and collaborate. The server includes two layers: (1) information repository, and (2) social networking.

The information repository layer includes two large-scale databases: (a) a virtual community databases that connects data from the local resource databases of suppliers providing physical information and human resources for use in disaster management (b) social media database which collects data from social media related to real-time emergent situations. The information repository layer will provide the knowledge foundation to the lower layer, the social networking layer.

The web-based social networking layer, on the other hand, provides information access interface, and communication and collaboration tools to users related to the emergency event within a suffering community, such as governmental agencies, relief workers, local residents, volunteers, and local disaster management organizations. The social network layer will provide multiple accessing interfaces including normal webpage, webpage for mobile devices, and text messaging to service needs of different users.

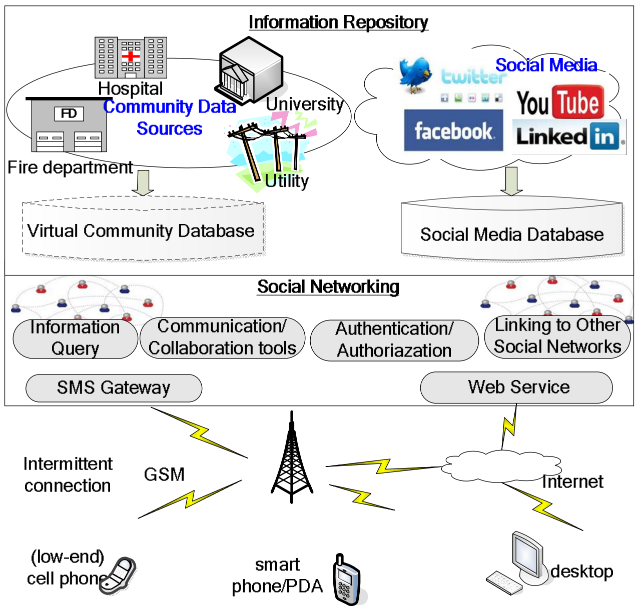

Figure 1: System architecture

The proposed disaster-management cloud provides a community-based, effective, and self-scalable cloud computing environment in which a diverse set of organizations and personnel can contribute their resources, such as data, knowledge, storage, and computing platform to the cloud. In this way, local communities, institutions/organizations, and individuals can seamlessly interact with each other to achieve massive collaboration in the emergency-affected areas. Semantic web technologies will be applied to facilitate integration of heterogeneous data from different originations and individuals. Moreover we recognize the widespread adoption of mobile handheld devices, such as smart phones, PDAs, iPads. Therefore, this project also will exploit the capacity of these mobile devices for disaster data collection, personal communication, alerts, and warnings by providing a set of Mobile Web-Service tools that can be used by search and rescue personnel via handheld devices.

3.2 Virtual Community Database

We use P2P architecture to manage distributed data sites of the effected community to form a P2P-based networked database. In this architecture, each organization maintains its own dataset and it also connects to one or more other organizations within the community. As a consequence, data sharing is performed on this community network. This P2P model allows a data site to easily join and leave the community network; it also allows owners of data sources to fully control access to and sharing of data without relying on a centralized server (a potential bottleneck). Moreover, the failure of one data site will not affect the functionality of the whole system. The aforementioned advantages make the proposed P2P architecture an ideal choice for disaster management systems in which data sharing must be set up quickly and easily with limited resources but the availability of a centralized server cannot be guaranteed.

To overcome data heterogeneity that originates from having different organizations involved in the community, we extend the distributed datasets with a semantic dimension. In particular, we utilize ontologies to unify the semantically related data in different sources. Namely, we defined a disaster-related ontology to facilitate data integration. Although ontology definition for disaster management is still very immature, there are some ontologies related to disaster/emergency [15-17]. We will adopt some of the existing ontologies. In addition, we will extend them with our new findings based on the knowledge inferred from the interviews with the relevant partners, papers, manuals and emergency proceedings. In particular, we will focus on some main subjects, namely, types of hazards and emergencies, meteorology, i.e., weather issues that might trigger an emergency situation, date and time, and most importantly, geographic concepts which can describe geographical regions affected by the disaster. Fig.2. illustrate snippets of disaster-related ontologies.

In the emergency-affected community, it is impractical to assume there is a global ontology defined for all participating data sources. Instead, we allow each data source to have its own ontology. Data sources can specify how the emergency-related concepts in their local ontology are mapped to concepts in disaster-related domain ontology. To enable sharing and integrating between participating datasets, ontologies of each dataset are extracted as a conceptual view over the data, which enable the access and query of the underlying data with ontology vocabularies. We take advantage the special property of disaster-related query, which focuses on concepts and relations about location, time, weather, and other disaster-related properties, to simplify the data integration.

To locate desirable data to construct mappings between dataset and to resolve queries efficiently, we propose to use a distributed hash tables (DHTs) [18]-based P2P network to implement a distributed ontology repository for storing, indexing and querying ontology knowledge. DHTs are a class of decentralized distributed systems. They provide efficient lookup services similar to a hash table. DHTs partition the ownership of a set of keys among the distributed participants. The indexing on the distributed repositories will speed up the searching process by only pushing down queries to information sources we can expect to contain an answer.

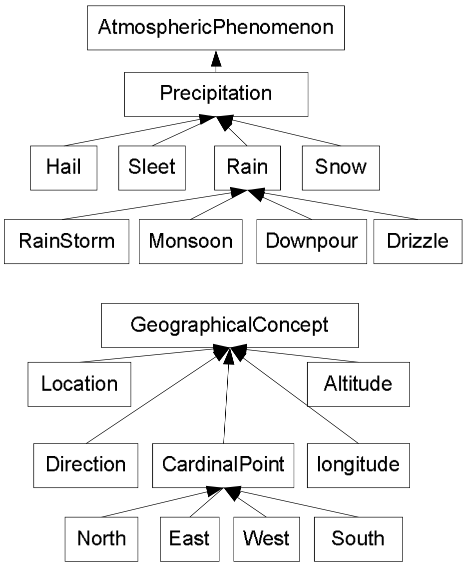

Figure 2: Snippets of disaster-related ontologies.

3.3 Virtual Social Media Database

Using social media, emergent information can be quickly disseminated out to the public. Moreover, in the case of a disaster, existing information infrastructure at the disaster location may break down and information may not be available through the usual channels. Social media may become the only platform for persons to report and acquire information about events and damages in their neighborhood, i.e., which is not available by other means. Furthermore, the state of the situation may change extremely quickly. For instance, hospitals might shut down abruptly due to electricity shortcuts. Information about such events are highly distributed among people, crews, and coordinating centers. Social networking sites such as Twitter are ideal media to update such status.

Hence to utilize the large amount of valuable information from social media to assist disaster management, we "crawl" data relating to disaster or emergency from multiple social media sources, i.e., Facebook, Twitter, MySpace, LinkedIn. This crawling is based on keywords extended with semantically-related words defined in the disaster-related ontology. We distinguish authoritative information from emergency related organizations, such as National Weather Service and Federal Emergency Management Agency at Facebook. Because disaster is time-sensitive, we collect information from social media at finer time granularities as if they are sensor streams.

All collected information is first converted to schema-less resource descriptor formats (RDF) to solve the problem of heterogeneous representation and to provide for maximum flexibility. To solve the problem of semantic heterogeneity, ontology is used to assist information integration, and enable the automation of discovery and composition. Besides the previously defined ontology, we define ontologies to model the special context property of social media related to disaster.

To improve the efficiency of search and query answering over the data collected from social media, ontological data is decomposed into atomic elements and then indexed MapReduce [19]. This parallel architecture is used not only for data indexing but also result integration. MapReduce provides the necessary simplicity and shields users from the details of parallelization, fault-tolerance, data distribution, and load balancing. Indexing will also be done at different granularities (e.g., at schema level, instance level) to fit different query needs and boost

3.4 Community Social Cloud Interface

The community virtual database and the social media database form an information repository for all parties in a disaster. We need to provide users with a good interface to access the information repository to give them a good situational awareness view and expand the capability of sharing information with emergency partners at all levels. Besides providing information, it is equally important to provide users with communication and collaboration tools to deal with disaster situations. For these purposes, we constructed a social networking site (SNS) that assists in each of the four phases of disaster and emergency management: mitigation, preparedness, response and recovery. We use SNS as the interface to our social cloud platform because SNSs have some special benefits compred with classical ICTs: SNSs can naturally assist the information flow between organizations by creating efficient inter-organizational networking [20]; SNSs can effectively increases the interaction among organizations [21]; SNSs allow the local community working as a whole unit to engage in overall higher levels of risk-taking [22], and solve problems of collective action more easily [23]. Furthermore, SNSs also facilitate the rapid dissemination of information and improve access to resources among network members.

Our social network focuses on the disaster community and opens to all parties related to the disaster. To avoid overwhelming users with the huge amount of information from the information repository and to maximize the accessibility of the information, in this task, we provide a flexible information query and browsing interface. Namely, we provide different query interfaces for different users and query scenarios. The system provides multiple interfaces including keyword-based, ontological graph-based, SQL-based, and Map-based, while supports unified access to the information repository. Ontologies defined from previous tasks are then used to standardize and extend the query.

Besides classical tools provided by general SNSs, such as status update, walls, wikis, blogs, we are working on implementing new tools for collaboration such as authenticated discussion board, document sharing tools (including group editing), group calendar, instant messaging, and web conferencing. The implemented collaborative tools will support the interactions between group members and support the group decision making process. Therefore, these tools will support team membership, roles and responsibilities.

We also implement different access interfaces for different users and scenarios. In particular, we use a web-based interface for users who like to access SNS from their desktop computers. We will also implement a special web interface customized for mobile devices such as smart phone/PDAs. Considering the possibility of intermittent network connection during the disaster and users equipped with only low-end phones, SMS-based service including notification, warning, and query is provided.

We also designed authentication and authorization mechanisms to give different access control to different organizations and individuals. This will identify the users with different roles to improve the trust of the social network. Other services based on the roles can be applied, for example, official can propagate warning information from the site, dialog channels can be provide between residents with officials.

4 Experiments

Extensive simulations have been performed on all aspects of the proposed system, such as query recall, precision, system scalability, overhead, etc. Namely, we created an experimental scenario to show that our semantics-based query does improve the overall search performance. We then use two Information Retrieval (IR) standards: precision and recall as the performance metrics. Precision is defined in Eq. 1 and measures the purity of the search results, or how well a search avoids returning results that are not relevant. The "document" in the IR definition represents a resource in our experiment. Recall refers to completeness of retrieval of relevant items, as per Eq. 2.

$$precision = \frac{|relevantDocuments \cap retrievedDocuments|}{|retrievedDocuments|}$$

(Eq.1)

$$recall = \frac{|relevantDocuments \cap retrievedDocuments|}{|relevantDocuments|}$$

(Eq.2)

To make the experiment easy to control, we simplify the ontology data, i.e., we use a small-sized vocabulary set to generate the ontology data; we fix the mapping relation to the *equivalentClass* relation and ignore all other mapping relations. We compare our ontology-based search with semantics-free exact-match-based search. For the exact-match-based search, a query only matches a keyword without caring about the keyword's specific meaning. For ontology-based searching, we vary the number of mapping neighbors (m) a node maintains from 1 to 3. For all the results returned, we compute the precision and recall. The results are shown in Table 1 and clearly indicate that our ontology-based search dramatically outperforms exact-match in both precision and recall.

Table 1: Performance comparison of semantics-based search and exact-match search

	Exact-Match	**Semantic (m=1)**	**Semantic (m=2)**	**Semantic (m=3)**
Precision	57%	100%	100%	100%
Recall	33%	64%	89%	97%

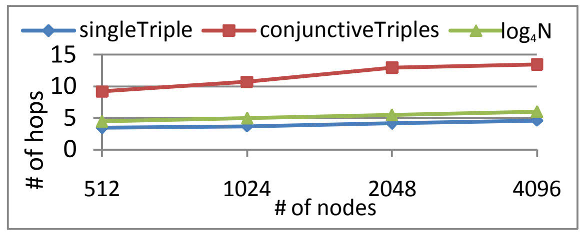

Figure 3: Query lookup efficiency

Fig. 3 presents the query lookup efficiency in the P2P-based virtual community database. It shows that our query evaluation (in particular, Conjunctive SPARQL query) based upon distributed hash table (DHT) indexing is efficient and scalable, i.e., scales by log(N) hops, where N is the number of datasets in the network and "triple" in the legend is the indexing unit. Meanwhile Fig. 4 also shows the screenshot of our community social cloud site aiming for collaborative management of red river flood in Fargo Moorhead area.

Figure 4: Screenshot of the RSC site

Overall, our preliminary results are promising. In addition, we are actually deploying our system through a website targeting the residents of Fargo-Moorhead. Here, the unique case of the annual Red River cresting will be used to assess the community response towards repeated possibilities of large-scale natural disasters. Namely, we will examine the overall users' experience levels with our system and use the feedback to improve our design, i.e., in terms of disaster warning, relief, and recovery. In particular, our evaluation efforts will use both quantitative and qualitative social science research methods.

5 Conclusions

In this paper we have proposed a cloud computing platform to maximally utilize all of the available information and human power from within and outside of a community to make information available to heterogeneous user groups. This enhances current information and communication systems to better manage natural and man-made disasters. Overall, our community-based cloud computing solution will allow for rapid scaling when needed and also provides significant flexibility and high cost reduction. Namely, cost reduction is driven by sharing resources and being more communal about computing, i.e., allowing the government's response to emergencies and disasters to be more agile.

Our community-based cloud computing solution retains data at the providers' sites, and thereby increases security as compared to the classical cloud model (in which data is stored in the third-party data center). However, there are still other security-related concerns which need to be considered further, e.g., access control, encryption, etc.

References

[1] D. Hilley, "Cloud Computing: A Taxonomy of Platform and Infrastructure-level Offerings", Technical Report, College of Computing, Georgia Institute of Technology, 2009.

[2] The National Institute of Standards and Technology http://www.nist.gov/index.html.

[3] Department of Homeland Security's Inspector General: http://www.dhs.gov/xoig/index.shtm.

[4] C. Loos, "E-health with mobile grids: The Akogrimoheart monitoring and emergency scenario", Akogrimo White Paper, 2006.

[5] C. Wattegama, "ICT for Disaster Management, Asia-Pacific Development Information Programme" e-Primers for the Information Economy, Society and Polity, 2007.

[6] F. Fiedrich, P. Burghardt, "Agent-based systems for disaster management," Communications of the ACM, Vol. 50, pp.41-42, 2007.

[7] C. Hagar, "The information and social needs of Cumbrian farmers during the UK 2001 foot and mouth disease outbreak and the role of information and communication technologies", The Socio-Cultural Impact of Foot and Mouth Disease in the UK in 2001: Experiences and Analyses (M. Döring, B. Nerlich editors), Manchester University Press. 2009.

[8] P. Currion, C. Silva, B. Walle, "Open source software for disaster management", Communications of the ACM - Emergency response information systems: emerging trends and technologies Vol. 50, No. 3, March 2007.

[9] O. Okolloh, O. Ushahidi, "Web 2.0 tools for crowdsourcing crisis information" (testimony), Participatory Learning and Action, 2009.

[10] M. Careem, et al, "Sahana: Overview of a disaster management system", IEEE International Conference on Information and Automation, 2006.

[11] SwiftRiver Website: http://swift.ushahidi.com/

[12] C. White, et al, "An online social network for emergency management", Proceedings of the 6th International ISCRAM Conference, May 2009.

[13] M. Turoff, S. Hiltz, "The Future of Professional Communities of Practice", WeB 2008 Proceedings (C. Weinhardt, S. Luckner editors), Springer, 2009.

[14] L. Plotnick, et al, "The Design of an Online Social Network Site for Emergency Management: A One Stop Shop", Proceedings of AMCIS 2009, 2009.

[15] W. Wang, C. Dong, P. Yang, "Ontology Modeling of Emergency Plan Systems", Fuzzy Systems and Knowledge Discovery, 2009.

[16] E. Little, G. Rogova, A. Boury-Brisset, "Theoretical Foundations of Threat Ontology (ThrO) for Data Fusion Applications", TR-2005 -269, November 2005.

[17] W. Xu, S. Zlatanova, "Ontologies for Disaster Management Response", Lecture Notes in Geoinformation and Cartography, 2007, pp. 185-200.

[18] Y. Li, Y. Feng, "Study of Data Search in DHT P2P Networks", Application Research of Computers, 2006.

[19] J. Dean, S. Ghemawat. "MapReduce: Simplified Data Processing on Large Clusters", Communications of the ACM, Vol. 51, No. 1, 2008, pp. 107-113.

[20] N. Kapucu, "Interorganizational coordination in dynamic context: networks in emergency response management", Connections, Vol. 26, 2005, pp. 35–50.

[21] J. Coleman, Foundations of Social Theory, Harvard University Press, Cambridge, MA, 1990.

[22] F. Fukuyama, Trust: The Social Virtues and the Creation of Prosperity, The Free Press, New York City, NY, 1995.

[23] E. Ostrom, Governing the Commons: The Evolution of Institutions for Collective Action, Cambridge University Press, Cambridge, UK, 1990.