

The RESCUE Disaster Portal for Disasters and Emergency Response

**Jay Lickfett, Naveen Ashish, Sharad Mehrotra
and Nalini Venkatasubramanian**
UC Irvine
{jlickfet, ashish, sharad, nalini}@ics.uci.edu

Jacob Green
City of Ontario Fire Department
jgreen@ci.ontario.ca.us

ABSTRACT

This paper describes the Disaster Portal, a multi-faceted information portal for use by citizens and emergency personnel during disasters and emergency response. This portal is the result of a collaboration between the RESCUE project at UC-Irvine and the City of Ontario, California Fire Department (OFD). The portal provides a wide range of real-time information in disaster situations, such as situation summaries, announcements, shelter information, and aggregated services such as family reunification and donation management. A first version of this portal was developed and deployed by the City of Ontario in September 2007. The current system serves as a base to develop and refine results from several areas of research which are being incorporated into the existing system to provide additional or advanced capabilities. We provide an overview of the portal's capabilities, followed by technical details on the system architecture and implementation. We describe the experience of deployment of the system in the recent California wildfires. Finally we describe work in progress on several advanced capabilities.

Keywords

Information portal, disaster response, information dissemination.

INTRODUCTION

The Disaster Portal¹ is an easily customizable web portal and a set of component applications used by first-responders to provide the public with real-time access to information related to disasters and emergency situations. Current features include a situation overview with maps, announcements and press information, emergency shelter status, and tools for family reunification and donation management. The Disaster Portal is being developed by the RESCUE Project², an NSF³ funded "large" information technology research project on advancing information technologies for disaster response. The goal is to make next-generation dissemination and situational awareness technologies available to first responders and the public, and to provide a platform for continued research by RESCUE project members. A pilot deployment of the Disaster Portal has been in use by the City of Ontario, California since mid-2007, and the city utilized the system extensively during wildfires in October, 2007.

PORTAL BENEFITS AND FEATURES

Current Applications

The Disaster Portal provides a variety of applications and capabilities that are useful to both citizens as well as emergency managers in disaster situations. These applications and services include:

SITUATION SUMMARY - The home page of the site displays summary statistics and/or most recent information updates for each of the applications deployed, as well as an interactive map-based overview.

ANNOUNCEMENTS / MAP ANNOTATIONS - Announcement messages provide information and instructions along with map annotations indicating the locations of current incidents.

¹ <http://www.disasterportal.org/ontario>

² <http://www.itr-rescue.org>

³ National Science Foundation

PRESS NOTIFICATIONS / PRESS RELEASES - Incident information and document attachments can be posted for media contacts, and notifications can be broadcast to press contacts via SMTP pager gateway or email.

SHELTER INFORMATION - A database of shelter facility information is used to manage and display current status, locations and directions, etc. pertaining to open emergency shelters.

FAMILY REUNIFICATION – A system for tracking displaced persons. Forms allow creation of "looking for you" information requests as well allowing a displaced person to provide contact information (i.e. from a kiosk set up in an emergency shelter).

DONATION MANAGEMENT - Allows users to post both requests and offers of donations in a categorized scheme, with a keyword search capability to assist in finding matches.

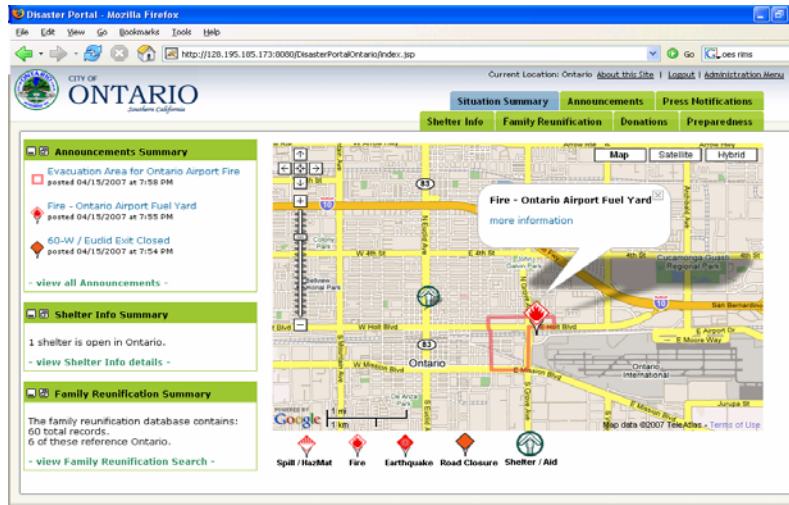


Fig 1. The Disaster Portal

The Disaster Portal dramatically improves communication between first-responder / government agencies and the public. Using the Disaster Portal, first-responders can quickly disseminate information to a wide audience. The Disaster Portal complements other systems such as telephone hotlines and can help reduce the load on those systems by allowing self-service to the information, and it also provides interactive maps and other rich content which cannot be easily shared via phone or other traditional means. Media contacts can be directed to access press releases on the site and subscribe to receive automatic email or pager notifications when new content is posted. The Disaster Portal also helps encourage civic participation by members of the affected community via the Donation Management and Family Reunification applications.

The Disaster Portal is easily customized to meet the different needs of the various organizations deploying it. Applications selected to be used in a particular deployment become available via tabs in the main navigation menu of the application, and it is also possible to add new content or functionality to the site. Furthermore Disaster Portal provides programming interfaces which allow it to be integrated with other IT systems. The Disaster Portal provides significant improvements in functionality and speed of development over ad-hoc web sites developed by individual organizations by providing a common set of capabilities customized for the disaster response domain.

DESIGN AND DEVELOPMENT

A schematic architecture of the Disaster Portal is provided in Fig 2 below. The design has been guided by several factors which include:

- The capabilities, described above that have to be provided by the portal.
- Interactions with and usage of the portal by citizens and first-responder administrative users.
- The sources of data that are ultimately used, including integration with other databases and information systems used by emergency response organizations.

- The ability to have integrated, multiple portal deployments to provide views of the situation at multiple levels in a geographical hierarchy (e.g. local, county, regional views.)

The portal is driven off information residing in the core portal database, which is a relational database containing information ranging from information about shelters, to announcements and press releases, to information that is dynamically collected off Web sites. Data is populated into this database in several ways. For instance, information about a newly activated emergency shelter or a press announcement can be entered by a portal administrator. On the other hand, real-time information such as information collected from online news stories or online sources providing information about public transport would be collected and entered to the database by automated sub-systems. Depending on the level of integration required these automated feeds may interact directly with the database or via higher level web services exposed by the Disaster Portal. Several information collection components are described in a later section. On the front end the portal is typically accessed through a web interface as illustrated in Fig 1. Users may also be provided information proactively, for instance through cell phone or pager alerts if they have requested these. Finally, we provide for portal administration and maintenance capabilities such as setting up of accounts etc. through web-based administration interfaces.

The implementation of the portal has been achieved by significantly leveraging open-source tools and software. This has primarily been driven by the cost efficiency and modularity and extensibility benefits that open-source technologies provide in general. This is important given that emergency and disaster response organizations operate under limitations regarding investment in IT and the use of open-source tools available free can alleviate such costs significantly. It is also highly desirable that any new IT components be such that they can be integrated with other IT components and tools that an emergency response organization may have already invested in. For instance many emergency departments at the city, county and state levels have invested in software such as WebEOC (WebEOC, 2007) and have expressed direct interest that new technologies such as the portal be integrated with packages such as WebEOC. Interoperability is thus another key aspect of the Disaster Portal. Finally scalability of technology deployment is highly desirable. In the specific context of the Disaster Portal it should be relatively easy and inexpensive to deploy new and multiple instantiations of the portal.

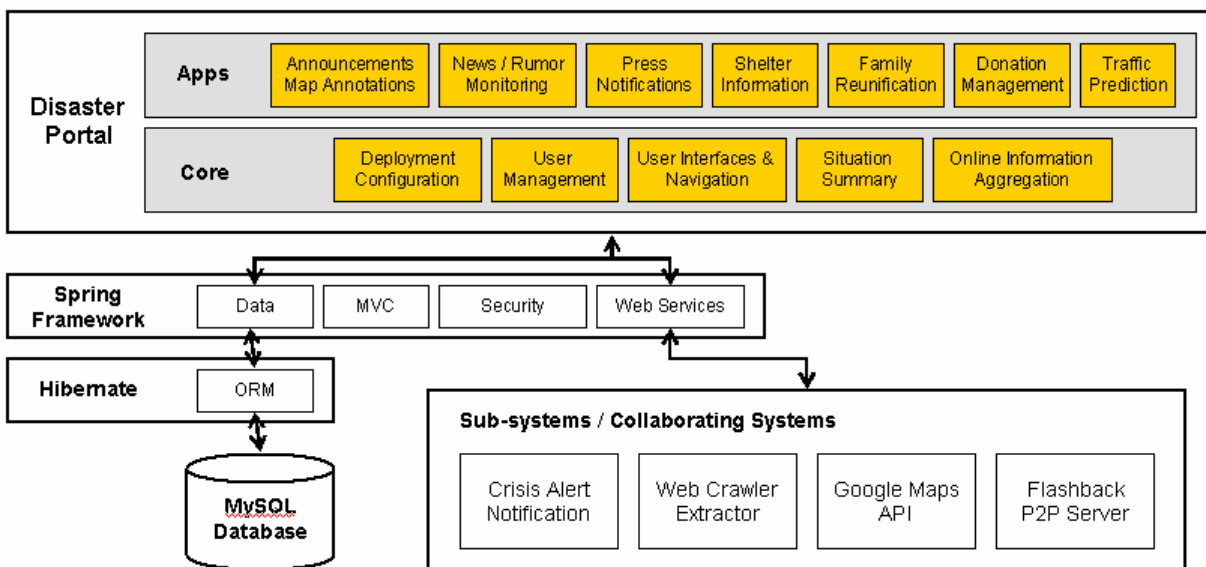


Fig 2. Disaster Portal System Architecture

The Disaster Portal application is implemented using Java technologies and utilizes the Spring Framework⁴. Spring is a Java application framework which provides a lightweight container for managing configuration and wiring of

⁴ <http://www.springframework.org>.

application components using dependency injection, as well as data and transaction management abstraction layers and AOP functionality. Spring's MVC, Security (formerly Acegi), and Web Services modules are used in conjunction with the Hibernate ORM to provide the foundation for the development the Disaster Portal.

The Disaster Portal application code consists of core functionality and a number of application modules which provide most of the functionality of interest to the end users. Core capabilities are those expected to be used by multiple application modules and include elements such as deployment configuration and user management. Online information aggregation capabilities and the representation of the present circumstances utilized by the 'Situation Summary' page are also grouped as core capabilities since they depend on multiple application modules. The Disaster Portal provides a set of standard JSP-based view templates for navigation and layout which are available for use by the application modules, with support for inclusion of dynamic maps provided via the Google Maps API through a custom JSP tag library included in the Disaster Portal core.

A number of application modules have been developed to provide end-user functionality described previously. These applications depend on the Disaster Portal core elements, and some also have additional integration with external sub-systems providing services such as web crawling, extraction (which we describe later), GIS services (e.g. Google Maps APIs), and dissemination (e.g. Crisis Alert System described below).

We currently deploy the Disaster Portal using Apache Tomcat 5.5 servlet container and MySQL as the relational database.

CASE STUDY: DEPLOYMENT IN SOUTHERN CALIFORNIA 2007 FIRES

The Disaster Portal has been available online to the public since September 2007. Its first substantial use by emergency managers for a disaster came during the fires in Southern California (in the Los Angeles, San Bernardino and San Diego areas) in late October 2007, when the Ontario Fire Department used it to provide information to their community. The Ontario area itself had several wildfires including the Kellogg, Eucalyptus and Walker fires. Of these Walker was the largest fire with 166 acres under threat and over 30 homes under evacuation.

During the Ontario fires, several elements of the Disaster Portal were utilized by personnel in the Ontario Emergency Operations Center established during the incident. The primary objective was to provide up-to-date information on the fires to the public and the media. The Disaster Portal was prominently featured on the City of Ontario homepage during the fires and public awareness of the site was encouraged via inclusion in reports made in traditional TV, radio, and newspaper media outlets. Over 20 announcements (many with map annotations) were posted, including announcements regarding location and extent of the fires, evacuation areas and instructions, shelter information, locations of downed power lines, and road closure information. Other information posted included how to obtain additional information (e.g. phone hotline numbers), and general cleanup and assistance information. An emergency shelter was opened during the fires in Ontario and updates about the shelter's status were posted on the Disaster Portal. The press notification application was also utilized to provide information directly to media contacts regarding mandatory evacuation orders which were issued for some areas during the fires.

Upon reviewing the web server logs we determined that approximately 5000 users (unique IPs) accessed the portal during the week of the fires, resulting in over 50000 page views. The Situation Summary and map page was most viewed with about 22000 views followed closely by announcements at 20000 views. The Press Notifications page had around 3000 views and Shelter Information about 700. Positive feedback was received from callers to the telephone hotline who found the information posted useful and in fact requested that more such information be posted online. Feedback from the city emergency management staff was also positive, and it is believed that the Disaster Portal did help in reducing the load on the 911 and telephone hotlines. This exercise also served to train several additional personnel at the OFD on the usage and administration of the Disaster Portal which will facilitate its use in the future and also resulted in additional feedback on the existing system.

In collaboration with the OFD, other cities, and contacts with volunteer organizations, a number of improvements and enhancements are now being identified for the portal. These include additional administration capabilities for emergency personnel and integration with existing emergency department information systems, such as WebEOC. Also planned are several new applications modules to provide additional information (for instance current status of public transport) in the portal.

ADVANCED FEATURES

The development of the Disaster Portal, to date and ongoing, is based on two primary considerations. We aim to provide practical applications and services of immediate utility to citizens and emergency managers, in parallel we also aim to significantly leverage many relevant pieces of IT research within RESCUE towards developing state-of-the-art portal technology. We provide a description of the advanced technologies that we are currently infusing into various components in the Disaster Portal. These include components for customizable alerting, family reunification, scalable load handling, unusual event detection and internet information monitoring.

Crisis Alert Dissemination Component – In addition to the current “pull” mode of information access (i.e., with users proactively accessing the portal) we also require a “push” mode where critical information can be sent out to certain potentially vulnerable or affected populations. A “Crisis Alert” component in the Disaster Portal automatically creates customized notification messages for a set of recipients who may be affected by a disaster or emergency situation based on administrator defined rules. These messages can be delivered via a variety of modalities including email, text messaging, and the RAPID peer-to-peer system also developed by RESCUE. This system can produce and rapidly disseminate both non-customized early warning messages to alert the population before a natural disaster occurs (earthquake, tsunami etc.) as well as post-disaster customized alert messages and updates in case of unpredictable and/or rapidly evolving events (fire, hazmat incident etc.).

The Crisis Alert system is essentially a policy-driven system which can be used for all kinds of hazard and disaster situations. The system manages a set of rules or crisis policies that summarize the knowledge about the various hazards and the vulnerabilities of the involved community. The rules provide specifications such as what action(s) should be taken in case of an event with certain properties, which message template should be used, and how the message needs to be customized given the current situation and intended recipients. A schematic overview of the Crisis Alert system is provided in Fig 3.

The event notification process is triggered by a *starting event* which can be a warning signal from sensors, another external system, or an input from an operator. In the first step we also determine what action should be taken given this triggering event. For instance an earthquake event of even moderate magnitude should call for an “evacuate” action. If such an action requires messages to be broadcast then an appropriate message template is chosen considering the event and action. The next step is to select the target recipients of the message which is done based on the affected area and the geographic information of the intended recipients. Individual(s) within the intended recipient organizations who will actually receive the message are determined. Finally the message is delivered through appropriate modalities, the current available delivery mechanisms include IP (PC and PDA clients), email, IP telephony and regular telephony.



Fig 3. Crisis Alert System Workflow

Family Reunification – The family reunification application provides the capability to post and search information

about missing persons and loved ones during a disaster. A planned addition to the Disaster Portal family reunification application is the ability to provide integration of crawling and/or searching of other missing person information sources on the web so that the user can effectively search many sites at once. This and related improvements will utilize results of ongoing research into issues such as crawling, information extraction, data uncertainty, data lineage, approximate query processing on text, and management of structured and unstructured data using the same infrastructure.

One of the research topics in this project is how to allow a user to find information about a missing person by doing keyword search, even if the representations of the person does not match the keywords in the query *exactly*. This feature of answering queries *approximately* is especially important due to the fact that information crawled from Web sites or typed in by people could have a lot of noise, and the user may have limited knowledge about the information of the missing person. For instance, when a user types in a keyword query "Michael Smith New Orleans", we want the system to be able to return entities with the following keywords:

- "Micheal Smyth New Oreleans"
- "Michael Smyth New Orelean"
- "Michael Smth New Oreleans"

These entities do not match the keywords in the query exactly, but they are returned because they could be of interest to the user. Supporting approximate keyword queries requires new efficient techniques on text search, especially for the case where we have a large amount of information in the data repository, and we want to answer queries very fast. We focus on developing algorithms that can answer a query efficiently without using a large amount of memory. Recently we have developed several such techniques. One of them is called VGRAM (Li et. al, 2007), which can improve performance (memory requirement and running time) of existing search algorithms by using short snippets of strings (called grams) with variable lengths. We also developed new algorithms (Li et.al, 2008) that can improve the speeds of existing algorithms by 5 to 10 times.

P2P Web Server for Scalable Load Handling – It is very likely that in cases of disasters and emergencies an information portal like the Disaster Portal (rather the Web server hosting it) experiences a sudden “spike” or “flash load” due to a very large number of users trying to simultaneously access the portal. To address this issue we have developed a system called Flashback (Deshpande et. al, 2007) which is an experimental web server which creates and utilizes a peer-to-peer infrastructure to address the problem of flash crowds overloading a traditional web server. The basic idea behind the Flashback technology is to transmit the flash load back to an end user browser, i.e., clients are recruited dynamically into sharing some of the server load. Such client involvement is done in a manner that is completely transparent to the end user i.e., without requiring any additional settings or software installation on the user end and also with guarantees of minimal impact (in terms of latency) to the user. Flashback is being integrated into the Disaster Portal to allow it to be deployed on typical web server hardware yet still remain effective during high-demand periods as might be expected during a disaster.

The architecture of Flashback defines it to be a self scalable system where the peer-to-peer networks grows and shrinks according to the load. The main contributions of Flashback can be summarized as:

- A fully functional and deployable peer-to-peer system for scalably distributing Web pages.
- A UDP based content dissemination protocol that is effective in situations of extreme *churn* (with clients entering and also leaving the peer network very rapidly).
- New and novel data structures for managing file *chunks*.
- Circumventing problem of peers being behind NAT devices (NAT stands for Network Address Traversal and NAT devices typically multiplex many end user nodes into the same IP address, an effect of this is that they do not allow new incoming connections).
- A technique, called “transported frame hack” which essentially allows the use of client as a server in a transparent manner, i.e., without any user intervention whatsoever.

Traffic Event Detection - This component provides the capability of detecting unusual activity or events in local traffic. It utilizes activity modeling (Ihler et al, 2007) in conjunction with live roadway loop sensor data from

CalTrans⁵ to provide information on current traffic patterns as well as predictions of near future conditions. Current efforts are being made to extend these models to track movements of populations in a given area. We are currently integrating this capability into the Disaster Portal where we can detect unusual activity or events in a region; these are determined based on data from sensors such as freeway traffic sensors, people counters etc.

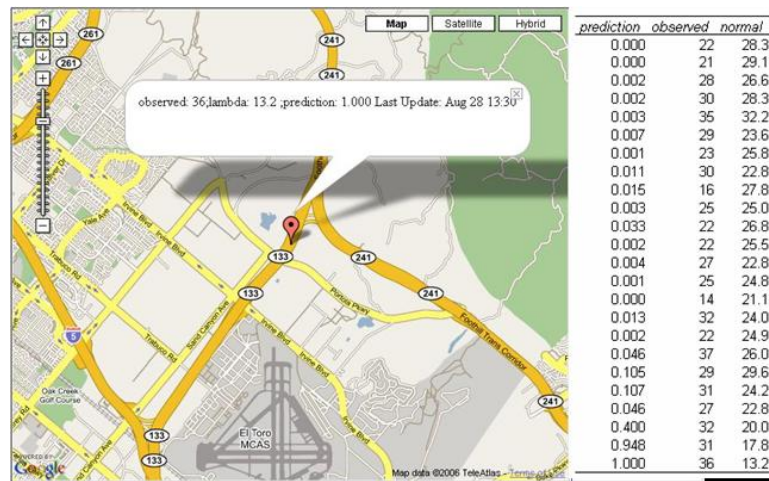


Fig 4. Unusual Event Detection

Our algorithm learns (unsupervised) normal traffic patterns and detects "event" activity defined as short bursty periods of unusually high or low traffic in real time using a time varying Poisson process model. The output of the inference algorithm is a probability that an event is taking place at any given time, along with an estimate of the extra number of people or cars or whatever you are counting that are attributed to that event. Also, we are creating software for a real time event detection system for freeway traffic in southern California. As an example of how this could be used in an EOC center, imagine evacuation planning during a bomb threat to downtown LA; the real-time event detection system is brought up, events at off-ramps in the past 3 hours is selected, and all off ramps near Dodgers stadium indicate the presence of an event, furthermore an extra 20,000 cars are predicted to have used those off ramps above the normal traffic; this gives a quick indication that something is going on in Dodgers stadium as well as an idea of the popularity of the event.

Internet Information Monitoring - As a disaster unfolds, a large amount of information related to the disaster quickly becomes available online. For instance we have news reports emerging at online news sources and also information posted in discussion boards, blogs, message boards etc. Much of this information is in text form and typically contains rich detail about particular events, new facts, messages, opinions etc. related to the disaster. Assimilating detail from such information or being able to query or analyze it is of interest from multiple perspectives. As an example, an important application that analysts at an EOC are interested in is that of "rumor monitoring", where emergency personnel are continuously monitoring news reports being published or broadcast about the disaster (including online news) to determine if any incorrect or false information is being reported about the disaster events and facts. Another application is that of damage assessment where risk and disaster management analysts look at published news and messages (amongst other information sources) to get an assessment of the damage to property that has occurred.

Currently such information assimilation and analysis is done manually and automated tools to do the same could significantly alleviate the time and resources expended currently. We have developed a module for the Disaster Portal for assimilation and monitoring of relevant online information. A schematic overview of this module is provided in Fig 5. The primary components are (i) A focused crawler which collects particular kinds of Web pages (for instance all news stories from local online news sources for the LA area that are related to an ongoing fire) at real-time, (ii) A "deep" extractor which is essentially a component that extracts detailed information (about events and facts) from the text in the Web pages, and (iii) A search and retrieval component that presents an interface to an end user.

⁵ California Department of Transportation

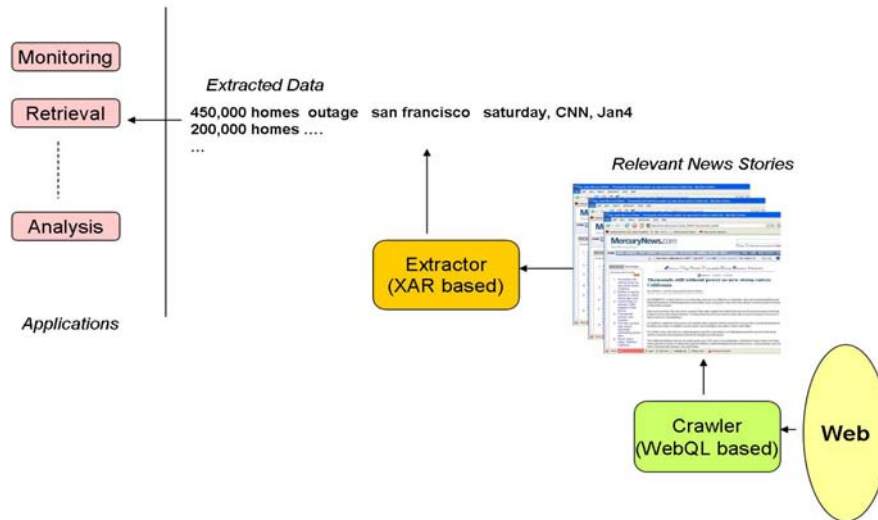


Fig 5. Internet Information Monitoring

The focused crawler takes as input specifications about the kinds of online sources to get data from as well as the kinds of Web pages that must be retrieved. For instance for local news monitoring we could provide a set of URLs for local online news sources, or pointers to sources such as Google News. We could also specify that only news stories and those related to a particular disaster (ongoing fire, landslides etc) should be retrieved. The focused crawler has been implemented using a commercially available web intelligence tool called WebQL (WebQL, 2007) which provides powerful capabilities for query and retrieving Web pages in a database (SQL) oriented fashion with control on the Web navigation as well. We have further augmented this with components for tasks such as classification and categorization of Web pages or links as they are retrieved or navigated at real-time.

The deep extractor extracts details about events or facts that are reported or described in the text. For instance a segment of text such as in Fig 5 contains details about a power outage event that is related to a disaster. As another example the text in Fig 5 contains details about injuries as a result of the disaster. An analyst may be interested in assimilating details about all instances of a particular event or fact, for instance all instances of power outages or all instances of injuries or deaths. The deep extractor component is able to automatically extract such details about particular (specified) events or facts from multiple Web pages at real-time. It has been implemented using the XAR extraction system (Ashish and Mehrotra, 2007), which is a state-of-the-art free text extraction system developed as part of the RESCUE project. XAR has powerful information extraction capabilities and has been designed to be intuitive and easy-to-use for new applications and users. It is open-source and available to the community.

Finally we provide a search and retrieval capability that provides appropriate browsing and search capabilities on such Web data over which we have performed deep extraction of details.

RELATED WORK

There are a number of initiatives in progress on building disaster information portals by organizations and individuals in academia, the industry, government, and non-profit consortia. While a comprehensive enumeration of these systems is beyond the scope of this paper, we provide a comparison of our work on the Disaster Portal with some closely related systems.

Sahana (Sahana, 2007) is an open-source system and effort for building disaster management applications. Sahana provides a suite of plug-and-play components using which one can build a variety of applications useful in disaster management and response. Key applications that have been built so far include missing persons registry for finding missing persons, an organization registry for relief organization coordination and tracking, request management system for tracking and matching specific aid or relief requests and offers, camp and shelter registry, volunteer management, inventory management to track supplies, and situational awareness. The Sahana project was initiated by a volunteer team of computer scientists in the aftermath of the 2004 Tsunami disaster. It has subsequently

received sponsorship from organizations like IBM and the NSF and has been used by governments in the Philippines, Indonesia and Sri Lanka.

At the application level there are significant similarities between the Disaster Portal and the Sahana system, especially considering the kinds of driving applications for the two systems. The key differentiation is that Sahana is primarily intended to be an open source, easy to use, application development platform addressing practical disaster management needs whereas in the Disaster Portal in addition to providing practical applications we are also developing and implementing state-of-the-art approaches to some of the particularly challenging problems. For instance the family reunification component in Disaster Portal (analogous to the missing persons registry in Sahana) is not just a set of entries for lookup as in Sahana (which is also useful) but additionally integrates information from a number of different online missing persons sources. Information matching (of things such as person names) is a challenging problem across such multiple sources, which we are addressing using disambiguation and string matching algorithms. Another example of an advanced feature in Disaster Portal is the integration of “Flashback” technology which essentially enables such a portal server to be responsible even under conditions of very high load, this is typical during disasters.

Another interesting system is Alertpedia (Alertpedia, 2007), originally called Alertearth. Alertearth started as a generic alerting system for immediately alerting citizens in case of disasters. Users could sign up for particular disasters of interest (say earthquakes) and such alerts would be issued over a variety of modalities, such as email, text messages, pagers etc. Alertearth has now expanded to Alertpedia and includes alerts and warnings for other hazards ranging from bio hazards to product recalls. The system continuously monitors hundreds of online and other sources to detect significant events and issue alerts. The focus of Alertpedia is exclusively on the alerting aspect which is one of the many aspects of the Disaster Portal. Our particular focus, as described in the Crisis Alert component, is on issues like action templates (as a consequence of the alert message), customization of messages, etc.

ONGOING WORK AND CONCLUSIONS

The major focus areas during the next phase of portal development will be 1) to expand the usage of the system to additional communities, 2) allow for multiple “cloned” instantiations of the portal at different locations and with possibly different agencies, and 3) incorporate results of additional research into either the core application or additional application modules.

With regards to 1) and 2) above, we are already in the midst of the some efforts to expand the scope of the portal usage, namely:

- After deployment at the City of Ontario, the Disaster Portal is now being installed and deployed at several cities in the Los Angeles and Orange County areas including the cities of Rancho Cucamonga, Upland and Aliso Viejo.
- We are in the process of coordinating with relief organizations such as the Red Cross that have expressed interest in the Family Reunification aspect of the portal.
- We are working with local community clinics and with members of the department responsible for public health information to explore the potential for mass alerting using the Disaster Portal.

The above expansion efforts have highlighted the requirement that the Disaster Portal architecture provide support for information interchange and integration across multiple agencies. Disasters that span multiple cities and jurisdictions are common and there may be a need for the Disaster Portal deployment at say Ontario to connect with and exchange information with the portal deployment at the City of Rancho Cucamonga. At a higher level, such as the county or state level, information may need to be aggregated from portal deployments at multiple cities to provide a situational “big picture”. We are in the process of developing capabilities in the Disaster Portal architecture that will facilitate seamless information interchange and exchange across multiple portal deployments.

We are exploring several options for facilitating utilization of the Disaster Portal software by additional cities, including licensing the system to a third party organization or by developing additional support material for deployment and training and releasing these along with the code under an open-source license. Cities or organizations interested in collaborating on the continued development of the Disaster Portal, or in running pilot deployments are encouraged to contact the authors for more information.

We would like to acknowledge the contribution of several members of the project “RESCUE” team that have contributed in some manner to the conceptualization and development of the Disaster Portal.

REFERENCES

1. Ashish, N. and Mehrotra, S. (2007) XAR: A System for Information Extraction from Free Text, *Technical Report UCI*
2. Li, C., Wang, B., and Yang, X. (2007) VGRAM: Improving Performance of Approximate Queries on String Collections Using Variable-Length Grams, *VLDB*
3. Li, C., Lu, J., and Lu, Y. (2008) Efficient Merging and Filtering Algorithms for Approximate String Searches, *to appear in ICDE 2008*.
4. WebEOC (2007), Whitepapers at <http://www.esi911.com>
5. Sahana (2007), <http://www.sahana.lk/>
6. Alertpedia (2007), <http://www.alertpedia.org>
7. Deshpande, M., Amit, A., Chang, M., Venkatasubramanian, N., and Mehrotra, S. (2007) Flashback: A Peer-to-Peer Webserver for Handling Flash Crowds *IEEE International Conference on Distributed Computing Systems (ICDCS)*.
8. Ihler, A., Hutchins, J., and Smyth, P. (2006) Adaptive event detection with time-varying poisson processes, *In KDD*
9. WebQL (2007), <http://www.ql2.com>