Improving Information Exchange and Coordination amongst Homeland Security Organizations

Terrance Goan
Stottler-Henke Associates, Inc.
1107 NE 45th St. Suite 310
Seattle, WA 98005
goan@stottlerhenke.com

Dr. Israel Mayk
C2 Directorate (C2D)
US Army Communications-Electronics
Research, Development and Engineering
Center (CERDEC)
Fort Monmouth, NJ 07703
israel.mayk@us.army.mil

Submitted to: 10th ICCRTS in the area of Homeland Security

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2005		**-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Improving Information Exchange and Coordination amongst Homeland Security Organizations	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **US Army Communications-Electronics Research, Development and Engineering Center,Fort Monmouth,NJ,07703**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**		**13**	

Improving Information Exchange and Coordination amongst Homeland Security Organizations

Terrance Goan
Stottler-Henke Associates, Inc.
1107 NE 45th St. Suite 310
Seattle, WA 98005
goan@stottlerhenke.com

Dr. Israel Mayk
C2 Directorate (C2D)
US Army Communications-Electronics Research, Development and Engineering Center (CERDEC)
Fort Monmouth, NJ 07703
israel.mayk@us.army.mil

Abstract

Command and Control for coordinated response to domestic terrorist attack will require the ability of federal, state, and local agencies to maintain awareness of the status, capabilities, requirements, response plans, and C2 procedures, etc. of the other collaborating organizations. While progress is being made in improving information sharing, the TOPOFF exercises have demonstrated that organizations still lack any substantial ability to coordinate responses to large scale events that involve dozens of local, state, and federal organizations. In this paper we describe progress made in the development of new information access services that provide for improved situation awareness. We have strived to develop a solution that enables User Defined Operational Picture (UDOP) functionality while respecting the unique information management practices of the collaborating Homeland Security organizations. Our system concept, Vista, employs an adaptive machine learning paradigm that supports a new form of context-sensitive information access, monitoring, and alerting that fills substantial gaps in existing Crisis Information Management System technologies. Experimental results demonstrate very substantial improvements in information access efficiency and provide strong evidence for the feasibility of the overall concept.

1 Introduction

Homeland Security (HLS) is rightly a high priority at all levels of the government, and huge investments are being made to train and provide basic resources to our country's first responders. We are also seeing dramatic changes in our intelligence agencies in order to improve information sharing across organizational boundaries (e.g., the formation of the Terrorist Threat Integration Center (TTIC) and the creation of information portals such as the Northwest Warning, Alert & Response Network (Northwest WARN)). Further, the DoD is making available critical assets through organizations such as the Army's Guardian Brigade which will provide military assistance to civilian law enforcement officials in emergency situations involving biological or chemical weapons.

Unfortunately, there are three very significant factors that complicate the situation. First, and most obvious, are the range of institutional and social barriers that have inhibited the exchange of information amongst federal, state, and local agencies. Intelligence sharing and coordination can be expected to be problematical due to the general lack of a unified and hierarchical command in situations involving dozens of collaborating governmental and non-governmental organizations which all operate with high degrees of autonomy and their own missions/goals (particularly in so-called "home rule" states).

The second complication is a general lack of system and semantic interoperability amongst HLS organizations. While modern information technology obviously has the potential to provide great improvements in situational awareness and coordination, in reality high technology is just now making substantial inroads into Emergency Management (EM). The DHS report following TOPOFF 2 highlighted this problem and in particular discussed the heavy reliance on the most basic forms of communication (e.g., face to face meetings, phone, fax, and radio). This problem together with the lack of a shared vocabulary obviously limited the ability of these heterogeneous EM organizations to produce a common operational picture.

Efforts are being made to improve the situation. For example, the Homeland Security Presidential Directive (HSPD)-5 directs the Department of Homeland Security (DHS) to develop and administer a National Incident Management System (NIMS) (see NIMS (2005)):

> *"This system will provide a consistent nationwide approach for Federal, State, and local governments to work effectively and efficiently together to prepare for, respond to, and recover from domestic incidents, regardless of cause, size, or complexity. To provide for interoperability and compatibility among Federal, State, and local capabilities, the NIMS will include a core set of concepts, principles, terminology, and technologies covering the incident command system; multiagency coordination systems; unified command; training; identification and management of resources (including systems for classifying types of resources); qualifications and certification; and the collection, tracking, and reporting of incident information and incident resources."*

Of course the practical benefits of these efforts may take substantial time to manifest. This leads to the third significant complication impacting the coordination of HLS organizations which is a simple lack of practice. The TOPOFF exercises are some of the few significant exercises involving a coordinated response to a large scale terrorist threat. The infrequency of these large scale exercises is not at all surprising given their high cost ($16 million in the case of TOPOFF 2). It is important to note that smaller scale exercises held at the state and local levels prove very effective in improving the ability of EM organizations to coordinate in crisis situations, but these very seldomly involve Federal agencies or National Guard participants. The resulting lack of familiarity, shared goals, and common understandings caused very substantial problems during the TOPOFF 2 exercise that will likely only be amplified in a real HLS crisis – causing problems in coordination as well as a simple lack of awareness of the full range of resources that could be brought to bear.

In this paper we describe progress made developing new technologies that can help overcome the above challenges and complement recent progress in the development and use of Crisis Information Management System technologies. Our overall goal is to improve situation awareness and thereby the ability to coordinate crisis response. More specifically, we have established a phased approach which first focuses on improving existing workflow processes used by EM organizations to access and monitor key information across institutional boundaries. Over the longer term we will develop methods for the efficient establishment and maintenance of a new level of semantic interoperability that will in turn allow a much higher level of automation in delivering the military's desired User Defined Operational Picture (UDOP) functionality that can provide improved awareness of unfolding events, the disposition of available resources, etc.

The remainder of this paper is organized as follows. First we review some of the key lessons we learned in our work with Washington State Emergency Management organizations who played key roles in the TOPOFF 2 exercise. We will then describe the current state of the art in

Crisis Information Management Systems and pinpoint the technology gaps we seek to fill. We will then describe the progress we have made in developing information access and monitoring tools that can streamline existing workflows and include some early experimental results. We will conclude with a description of our future efforts.

2 The Practicalities of All Hazards Emergency Management

In order to understand how the work described herein can support HLS operations, it is first important to place HLS in the context of common EM practice. There are in fact several key issues that must be considered in order to successfully field technologies that can support HLS. First, the procedures followed in HLS operations are not separable from general EM practice. Instead Emergency Operations Centers (EOCs) operate under so-called "all-hazards" emergency response strategic plans which stipulate under what conditions the EOC is activated, who ultimately directs the response, and how outside resources (e.g., military personnel) are acquired and exploited. More importantly these plans specify that basic policies, procedures, and chain of command are the same whether an emergency is due to an earthquake, a disease outbreak, or a terrorist attack.

Second, consider that in "home-rule" states such as Washington State, a county EOC must coordinate dozens of largely autonomous government and non-government organizations – each with their own constituencies, resources, and needs. Even in non-home-rule states, coordination with autonomous non-governmental organizations poses the same challenges to command and control since information sharing and general cooperation is voluntary. For these reasons communications are ad hoc relative to that witnessed in military organizations.

Finally, we must consider that emergency response coordination, while the point of emphasis for the DHS and the press, represents a very small portion of the responsibilities of EM organizations. The full time staff of these organizations spends the vast majority of their time in tasks related to emergency preparation and recovery. Further, since EOC's are seldom fully activated and training exercises are infrequent, one can expect personnel to face significant challenges in exploiting technological tools when a crisis does occur. These facts highlight how critical it is that tools be easy to use. Perhaps even more importantly, technologies that solely support the least likely of events (e.g., large scale HLS events) are unlikely to be adopted by EM organizations.

2.1 The Joint Operations Center Concept

Another important concept is that of the Joint Operations Center (JOC), which according to state and local EM strategic plans, will be established by the local FBI office in the event of a significant HLS event. The purpose of the JOC is to provide a location where trusted representatives of Federal, State, and Local agencies can meet face to face to resolve critical issues and support a well coordinated response.

From the perspective of Seattle and King County EM, the JOC serves little purpose since the liaisons sent to the JOC will most often not be key decision makers, nor does it replace anyone's executive, political, command, or legal authority. Further, as we will discuss in the next section, it appears that modern communications technology provides more than sufficient support for executive decision makers to coordinate from a distance. This finding is in line with the hypotheses presented by Mayk and Klose (2004) with regards to military C2.

The JOC in TOPOFF 2 was the source of several problems according to Seattle's then Director of Emergency Management, Jim Mullen:

> *"During the exercise, an unauthorized evacuation was in the early stages of planning before our EOC—the command and control center for the event—even heard about it. Fortunately, we were able to stop it, but the lesson was that if you have a place where people with initiative but without authority are congregated, they will sometimes do more harm than good"*

3 State of the Art in Crisis Information Management Tools

In order to support the activities of the EOC, each is equipped with a wide range of communications equipment including FEMA National Warning System (NAWAS) terminals, Emergency Alert System (EAS) radio, Computer Aided Dispatch (CAD) systems, lines to the ATF and FBI, and digital and analog phone systems. These systems provide the baseline means for maintaining situation awareness and for supporting coordination with outside organizations.

EOC's are also undergoing some dramatic changes with regards to technology support. In recent years, teleconferencing systems have been installed, and EOC's are making a gradual shift from a reliance on grease pencils and laminated boards/maps to modern crisis management tools. For instance, the three Washington State EM organizations examined in our research use a tool called WebEOC to support the management and online dissemination of crisis information. This tool offers a variety of so-called "status boards" that are used to post and track the happenings within the collaborating organizations. WebEOC and similar tools such as E Team at their core are simple information sharing tools – essentially replicating physical bulletin boards that have been used in EOC's for many years and making them accessible over computer networks. Other features allow for incident replay for training purposes, alert dissemination, resource tracking, and limited incident/resource mapping.

Despite the promise of these tools, there are a number of emerging difficulties. First of all, according to a recent Department of Justice study (DOJ, 2002) there are approximately a dozen different Crisis Information Management Systems (CIMS) in use by HLS relevant organizations throughout the country, and the developers have come late to the concept of interoperability. Progress is being made to develop an Emergency Management XML schema to improve interoperability, but as yet this schema is limited to the so-called Common Alerting Protocol (CAP) which was recently released by a consortium of CIMS developers (CAP, 2004) and provides a template for emergency alerts. Because of its limited nature, CAP does little to support advanced incident coordination applications.

Another challenge facing users of existing tools relates to the supported information access paradigm. While CIMS tools allow for inter-organizational publishing and sharing of information across computer networks, information publishers and consumers are often out of synch. That is, the information "pushed" to consumers is often not relevant to their ongoing tasks and the information "pull" mechanisms may require much manual searching and monitoring. In fact, the EM organizations we examined in this research had to assign teams of personnel to continually monitor various WebEOC status boards and external Web data sources (e.g., news coverage, road closures, etc.) during activations in order to ensure key information is delivered to those that need it. Of course the challenges of timely information access and fusions will become even greater as the variety of information and information providers increases to allow monitoring of available hospital beds, power outage areas, the course of chemical plumes, etc.

4 Opportunities to Improve HLS Emergency Response

TOPOFF 2 and the experiences of state and local EM organizations in more conventional operations suggest that there remain many opportunities to improve HLS situation awareness and response coordination capabilities. Obviously, communications and coordination capabilities amongst HLS-relevant DoD, non-DoD, Federal, state, and local organizations is not what it should be. In the event of significant domestic terrorism, response time is critical and missteps can have extremely high costs in human lives. Therefore, new technologies are required that can dynamicallyexploit the output of new information providers to offer both vastly improved information/situation awareness and the ability to coordinate crisis response. With that said, we recognize that new technologies face a number of challenges in achieving adoption by EM organizations – most substantially the need to fit with existing workflows and systems and provide immediate benefit in day to day operations, as well as maintain low training requirements. With this in mind we are pursuing a phased R&D approach where early software deployments can inform us as to the remaining areas of need.

We currently are pursing three key opportunities, with the first being our initial focus.

1. **Intelligent context-sensitive information access, monitoring, and alerting across organizational boundaries.**

   Seizing this technological opportunity can provide context relevant information access so that EOC personnel, military commanders, intelligence analysts, etc., can quickly identify and track the information that they know they need to make time-critical decisions, as well as provide alerts with decision impacting information they did not even recognize they needed.

   This new information access capability exploits a model of the current operational context to provide automated Web site and status panel monitoring, federated search across cooperating organizations, and to provide users with an ongoing awareness of the information and knowledge being generated and exploited across HLS organizations that fills local information needs.

2. **The efficient establishment and maintenance of semantic interoperability.**

   Over the longer term we are seeking to develop techniques that can build upon the progress made by US Army C2 CERDEC in the development of C2XML (see Mayk and Klose, (2004)), the EM-XML consortium in the development of CAP (CAP, 2004), and the DHS's recent efforts to develop the National Response Plan (NRP) and the National Incident Management System (NIMS) in order to produce a mediating ontology that will be necessary to achieve interoperability across the legacy and future systems operated by HLS relevant Federal, DoD, and local EM organizations.

   This type of rigorous information model has much to offer, including the potential of allowing state and local EM organizations to bring to bear the valuable collaborative planning systems (e.g., CAPES/MC2) employed by the US military. It will also open the possibility of lower-cost, multi-agency, computer based training exercises. This type of asset sharing represents a novel and valuable direction in HLS.

3. **A vastly improved capability to provide awareness of unfolding events, the disposition of available resources, etc., through a User Defined Operational Picture.**

The aforementioned ontological foundation may provide a sufficient basis for the automated processing and formalization of C2 information such as significant event messages, EM tasking orders, Operation Orders (OPORDs), etc., and to support information system integration. These new capabilities will allow us to move beyond the current generation of manually updated electronic bulletin boards, and on to the production of a common operating picture that can be tailored to the user organization.

Not only will such displays offer the EM community a substantial improvement in situational awareness, it will prove an essential capability in achieving effective coordination with organizations outside the local civilian EM community.

5 The Vista Concept

Our initial *Vista* prototype focused on the task of extending and evaluating an existing context-sensitive information retrieval and monitoring technology (see Aware (2005)) in order to improve the ability of users to locate and maintain awareness of operation critical information. In particular we have prototyped and tested a new search scheme that mitigates the inevitable problem of a users' query terminology not matching with that employed by the authors of the desired information. This problem crops up in all information retrieval tasks simply because it is difficult for the information consumer to fully anticipate the terminology the information producer has utilized, and just as commonly users' queries can unintentionally match information that is irrelevant to their task. These challenges are particularly obvious in situations where information is being accessed across organizational boundaries, such as in the case of EM/HLS operations.

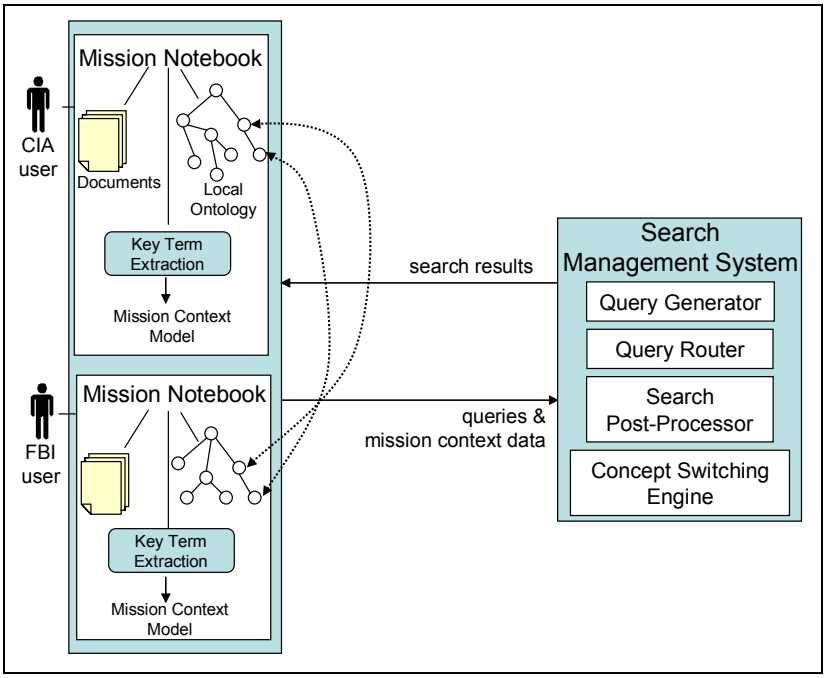

Figure 1. Vista Prototype Overview

The scheme we have implemented has the following primary components.

1. **Context Modeling.** Our search technology has, at its core, the ability to derive a model of the user's information need through an analysis of context information. This context

typically is composed of a set of rated (relevant / non-relevant) documents, but can also be derived through much less explicit means by monitoring the user's information management operations. In the context of this project, this may involve the status of personal or EOC activity logs complete with hyperlinks to information used to form decisions.

2. **Multi-search.** Leveraging the user's operational context, our technology is able to produce multiple refinements of user queries in order to scan a topic area. In the current effort we have made a number of improvements to this technology that make it feasible for *Vista* to autonomously search for information relevant to the user's expressed (via the operation context) information requirements. This advance required new techniques that are able to distinguish key terminology (what we call "query anchors") from other terminology that is only vaguely relevant (see **Figure 2**).

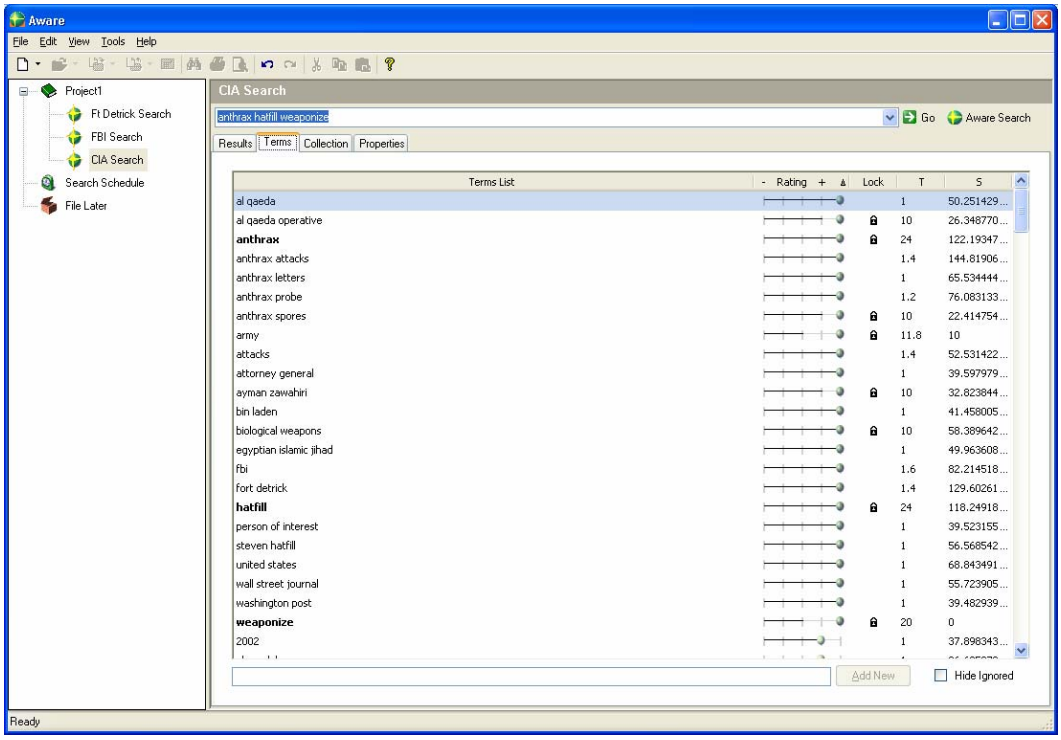

Figure 2. Vista's Mission Context Model. Vista processes the documents contained in the user's Mission Context (i.e., the document collection) and extracts out the key terminology defining the user's requirements. This list of terminology is continually refined by the user's rating of documents and through Vista's concept switching functionality. Note that terms have a user modifiable rating, a Traction score (T), and a system generated score (S).

3. **Traction.** One of our biggest successes in this project has been the prototyping of a method that is able to measure the "traction" of query terms in delivering information that the user requires. Based on binary relevance feedback, *Vista* is able to quickly identify which query anchors are most relevant to the user's current tasks, and most effective at identifying the required information (see **Figure 2**).

4. **Enhanced concept switching**. The other major advance we have made in our prototyping efforts is an extension of so-called "concept switching" (see Schatz (2002)). Concept switching is a mechanism for dealing with situations where the information seeker may not know the appropriate query terminology needed to access information

developed in another community/organization. Concept switching starts with the one core concept in the user's context, generates a set of related concepts in that space, and then uses this broader set to locate the most likely matching concept in the target community. In our prototype implementation of an enhancement of this concept, we allow the automatic acquisition of terminology by a user's (or community's) *Vista* system from the terminology (ontology) employed in another community's *Vista* system, based on the presence of shared interests. We then employ the aforementioned *traction* heuristics to validate that the terminology acquired is indeed relevant to the user's task.

5. **Context-based information evaluation and filtering.** Finally, *Vista* is able to evaluate the relative quality of discovered information based, not only on a limited query, but on the entirety of the search context. This more comprehensive model leads to greater accuracy in information ranking/filtering and even enables the system to measure the novelty of new information (which may be indicative of a significant turn of events).

6 Preliminary Experiments

In order to test the efficacy of our newly prototyped technologies, we designed two experiments. In the first experiment, we used human test subjects to compare the efficacy of search via traditional search tools versus the Vista context-aware search methods. In the second we tested the viability of our enhanced concept switching method.

In these preliminary experiments we relied on a relatively simple search task specification in order to evaluate the effectiveness of our techniques. In this test scenario we have two communities of users, the CIA and the FBI, who are investigating the 2001 Anthrax letter attacks. Each community has its own unique search objectives (the CIA is investigating links to international terrorist organizations and foreign states like Iraq, while the FBI is examining domestic suspects) and terminology, but clearly there is a substantial overlap in information that would be deemed relevant to their respective investigation.

6.1 Testing Context Based Search

For this study we recruited ten subjects from within our company, whose occupations ranged from technical writer to software engineer. One half of this group was randomly selected to act as a Control group, and were asked to produce a total of four queries to the Google search engine and review at most ten pages per query. We did not prevent the users from examining multiple pages of search results, but each page of ten results counted toward their four query limit. Subjects were then asked to rank the results as Unrelated, Related, or Highly Related.

The other half of our study subjects acted as the Experimental group. Each person in this group was only allowed two queries, and was given the same instructions, but was also asked to provide our search tool with the pages they found highly related to the topic. Our search tool used the resulting set of texts as a context for one additional, automatically constructed query which was used to fetch fifty results and rank them – reserving only the top 20.

The results of this experiment were extremely encouraging. We found that our search system was able to make very good use of the small efforts taken by the Experimental group to provide feedback. In fact, it was able to locate results that were considerably better than those located using the queries constructed by hand by the control group after they had analyzed two pages of search results. In particular, while the Control group saw a 56% increase in the number of high quality pages resulting from their second two queries (in comparison to the first two), the

Experimental group saw a jump of 127% (see **Table 1**). The results were equally impressive when considering documents rated as either Related or Highly Related (see **Table 2**).

Also note that we separated Test Subject 10 from the rest of the Experimental group, because he had some amount of training in the use of the software and was therefore better able to provide texts that would direct the system. It is interesting that Test Subject 10 did not do substantially better than those subjects who had not been given any instruction in how to best train the system.

Control Group	**Highly Related (1st 2 queries)**	**Highly Related (2nd 2 queries)**	**% change 1st half to 2nd half**
Subject 1	1	3	200
Subject 2	5	9	80
Subject 3	1	3	200
Subject 4	2	0	-100
Subject 5	3	0	-100
Average	**2.4**	**3**	**56**
Experimental Group			
Subject 6	4	10	150
Subject 7	2	6	200
Subject 8	4	6	50
Subject 9	5	9	80
Average	**3.75**	**7.75**	**120**
Subject 10	**7**	**16**	**128.6**

Table 1. The Search for Highly Related Pages

Control Group	**Total Related (1st 2 queries)**	**Total Related (2nd 2 queries)**	**% change 1st half to 2nd half**
Subject 1	12	5	-58.3
Subject 2	9	14	55.6
Subject 3	4	7	75.0
Subject 4	4	1	-75.0
Subject 5	5	5	0.0
Average	**6.8**	**6.4**	**-0.6**
Experimental Group			
Subject 6	9	14	55.6
Subject 7	5	18	260.0
Subject 8	15	12	-20.0
Subject 9	9	14	55.6
Average	**9.5**	**14.5**	**87.8**
Subject 10	**11**	**20**	**81.8**

Table 2. The Search for Related Pages (including Highly Related)

6.2 Testing Enhanced Concept Switching

In this preliminary experiment we utilized a simpler evaluation scheme for expedience. In this experimental set up we constructed two search profiles – one representing the FBI and one the CIA. In each case we had previously conducted intensive manual searches to identify pages relevant to the search profiles (as well as a set of off-topic pages). We then began a batch search process that proceeded as follows:

1. An initial query was executed for the CIA profile (e.g., 'anthrax terrorist iraq'). Any documents that were returned by the search engine were compared to the associated list of relevant and off-topic documents and Vista was provided with the appropriate feedback as if the human user were involved.
2. An initial query was executed for the FBI profile (e.g., '"anthrax letters" "ames strain"'). Feedback was provided as above.
3. If any document was located that was judged as relevant to both the FBI and the CIA, Vista allowed terminology to cross from one profile to the other. In particular, the query terms that matched the particular results were tentatively assumed to be relevant to the other search profile. Following the crossover, traction was employed to validate the utility of the term in the search profile – degrading it if the term failed to produce more relevant results.
4. A series of interleaved queries (CIA, FBI, CIA, etc.) was then automatically produced, each followed by appropriate feedback and terminology crossover operations.
5. Finally, the number of relevant documents retrieved was tabulated

In this experiment we allowed a total of 40 searches per search profile (simulating multiple users in each organization.)

The results from this experiment were surprisingly good. The test runs employing our enhanced concept-switching returned an average of 45% of the known good FBI documents (as compared to only 32% when sharing was not employed), and an average of 35% of the known good CIA documents (as compared to 31% when sharing was not employed). These are substantial improvements in recall for a technique only a few months under development.

Perhaps even more interesting is the quality of the terms that participated in the terminology cross-over (e.g., "Al Qaeda operative," "anthrax attacks" "Ames strain of Anthrax"). These terms are surprisingly good given the remarkably simple statistical method we use to identify these multi-word terms. This success is a strong indicator that our traction scheme is very successful, since terms extracted from documents are continuously reevaluated based on their ability to produce good results – the more good results a term generates, the more likely it is to be shared.

7 Related Work

7.1 Crisis Management

A number of companies are now offering crisis management software which supports the critical tasks of EM. These tools support information sharing in applications such as homeland security, business continuity, disaster preparedness and recovery, event management, and training exercises (see a survey by the Department of Justice (DOJ, 2002). There also exist a number of examples of crisis decision support concepts under development within the research

community. One example is the EMERRS (Emergency Regional Response System) application (Pohl, 2001) which seeks to promote effective emergency response planning in urban response units. EMERRS is intended to integrate data from disparate sources into a single coherent view. But while the investigators intend for this system to provide a disciplined decision-making environment, it is currently limited to data integration capabilities and does not support UDOP functionality.

7.2 Ontology Alignment and Information Access

Tools for merging or aligning ontologies help users find similarities and differences between source ontologies. These tools either identify potential correspondences automatically, or provide the environment for the users to find and define these correspondences, or both. Merging/Aligning tools are often extensions of development tools. Prominent tools include: PROMPT (Noy and Musen, 2000), ONION (Mitra, et al., 2000), Chimaera (McGuinness et al., 2000), and FCA-Merge (Stumme and Madche, 2001). See (Noy and Musen, 2002) for a survey. While the approaches employed by these different systems vary substantially, one thing that all automated methods have in common is that they rely on the structure of the ontologies to drive the alignment process. This differs substantially from the approach used by Vista, where we seek to exploit the text matching different information extraction ontologies to aid in proper alignment and merging.

Schatz (2002) describes their Interspace Prototype, an analysis environment supporting semantic indexing on community repositories. The most significant aspect of this work is their approach to concept extraction and concept spaces, which use semantic indexing to facilitate concept navigation. Of particular interest is their notion of "concept-switching" as a mechanism for dealing with situations where the information seeker may not know the appropriate query terminology to use in another community's knowledge space. Concept switching starts with one concept in one community space, identifies a set of related concepts in that space, then locates similar related sets in another space. Simulated annealing (or the like) can then be used to identify the closest match to the original concept.

8 Conclusions and Future Work

We have established the technical feasibility of our Vista concept in multiple ways. In this effort, we have worked interactively with decision makers in three Washington State EOC's to establish a good fit between technology and need – with a focus on how best to facilitate information sharing and coordination. Second, our experimental results (although limited) provide evidence that our techniques for context-aware information access, monitoring, and alerting will provide immediate benefits to the EM/HLS community, thus promoting user adoption.

Our Vista project is a work in progress. Most recently, we commenced another round of interviews and workflow studies with Washington State EM organizations. Based on the accumulated information, we have developed a preliminary design for the first version of *Vista* that will be deployed in an EOC. This refined concept will enhance the common operating picture for Emergency Management workers during the emergency response phase by:

- Automating manual searches for critical information on agency internet and intranet websites
- Providing a combined and organized view of the incident and its various aspects through intelligent displays and inference

- Finding underreported pieces of relevant data and routing that information to the appropriate decision-maker
- Autonomously monitoring external and internal electronic data sources and integrating changes into the common operating picture

During the other phases of Emergency Management, Vista will continue to monitor information sources and alert users when triggered by specified events.

The potential benefits of Vista's improved situational awareness is achieved significant buy-in from Washington State EM/HLS organizations and they are actively involved in our requirements analysis and Beta software testing.

9 Acknowledgements

This work was supported by US Army CECOM under the Small Business Innovative Research (SBIR) contract W15P7T-04-C-K610.

10 References

Aware (2005) http://www.awaresearch.com

CAP (2000) CAP 1.0 - OASIS Standard. http://www.incident.com/cap/docs.html

DOJ (2002) Crisis Information Management Software (CIMS) Feature Comparison Report. http://www.ojp.usdoj.gov/nij/pubs-sum/197065.htm

Mayk and Klose (2004) Experimenting with C2 Applications and Federated Infrastructures for Integrated Full-Spectrum Operational Environments in Support of Collaborative Planning and Interoperable Execution, XBML/XMSF Workshop on Applicability of the Command & Control Information Exchange Data Model (C2IEDM) for Data Management and Data Interoperability.

McGuinness et al. (2000) An environment for merging and testing large ontologies. In A. G. Cohn, F. Giunchiglia, and B. Selman, editors, Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR2000). Morgan Kaufmann Publishers, San Francisco, CA, 2000.

NIMS (2005) http://www.nimsonline.com

Noy and Musen (2000) PROMPT: Algorithm and tool for automated ontology merging and alignment. In Seventeenth National Conference on Artificial Intelligence (AAAI-2000), Austin, TX, 2000.

Noy and Musen (2002) Evaluating Ontology-Mapping Tools: Requirements and Experience. Workshop on Evaluation of Ontology Tools at EKAW'02 (EON2002). 2002.

Pohl (2001) “Integrated Decision-Support Capability for Enhancing Crisis Management and Increasing Response” In the Internet Journal of the National Institute for Urban Search and Rescue (NIUSR). Vol. 1.

Schatz (2002) Concepts across the Interspace: Information Infrastructure for Community Knowledge, IEEE Computer, January 2002 Trends issue.

Stumme and Madche (2001) FCA-Merge: Bottom-up merging of ontologies. In 7th Intl. Conf. on Artificial Intelligence (IJCAI ’01), pages 225–230, Seattle, WA, 2001.