

**Fifth-Year RESCUE Annual Report –
10/01/07 to 09/30/08**

Responding to Crisis and Unexpected Events

ITR

Collaborative Research: Responding to the Unexpected

National Science Foundation Award Numbers:

IIS-0331707, University of California, Irvine

IIS-0331690, University of California, San Diego

1 OVERALL PROGRESS IN YEAR 5 – EXECUTIVE SUMMARY

RESCUE - *Responding to Crises and Unexpected Events* – is a five-year research project funded under NSF’s Large Information Technology program. This project involves seven major research institutions located throughout the U.S.; over 40 faculty and senior research staff members with backgrounds in computer science, social science, and engineering; over 60 graduate and undergraduate students; and 30 government and industry partners. This report provides a summary of progress and achievements in Year 5.

1.1 Year 5 Progress

Year 5 of the RESCUE project has been filled with many significant achievements. As a result of a refocusing of research priorities which occurred in Year 3 – prompted by a set of recommendations by NSF and its site review team – RESCUE directed its research to achieve six overarching strategies. Progress under each of these strategies is discussed below:

1 Structure RESCUE research to focus on a small set of problem-focused, multidisciplinary research projects that are driven by end-user needs and that also offer significant opportunities for scientific contributions.

[Progress in Yr 5: Measurable achievements have been made in all six major project areas: 1) Situational Awareness from Multimodal Input; 2) Robust Networking and Information Collection; 3) Policy-Driven Information Sharing; 4) Customized Dissemination in the Large; 5) Privacy Implications of Technology; 6) MetaSIM. For each project, a government or industry partner was named as part of the research team in order to drive research towards end-user needs. Where possible, research results were tested by government partners in actual drills or exercises, or during real crisis events, e.g., the southern California wildfires. While each project had individual research nuggets, the major contribution to scientific research was the comprehensive and multidisciplinary nature of the entire research program. Arguably, the depth at which specific problems were investigated demonstrated that information technology solutions do offer real opportunities to significantly improve crisis response in the U.S. The research team has been prolific in its presentation of research findings with over 500 publications so far. In addition, the validation of our research in the context of real-world problems is demonstrated with the great number of government and industry partners who have actively participated in the RESCUE project.]

2 Create a set of living laboratories and simulation tools that serve as testbeds which mimic “real-world” conditions for regional and incident-level crises and that reflect RESCUE’s mission and objectives.

[Progress in Yr 5: This has been a real strength of the RESCUE project. In total, the project has developed or participated in four major testbeds: the transportation testbed that includes a major portion of the southern California highway network; the UCI Responsphere testbed that creates a campus-level pervasive environment that supports a variety of networking and sensing capabilities; the San Diego GLQ/Extreme Networking testbed that served as a mobile living laboratory for deploying and testing a variety of communication technologies in various locations and drill sites, including exercises with the UCSD Campus, San Diego Police and Regional Metropolitan Medical Strike Team; and the Champaign, Illinois testbed where local emergency response officials are actively involved in evaluating information sharing technologies as used in simulated disaster scenarios. The RESCUE program has been fortunate in that for each testbed, high-level government officials have been active participants in both the research program as well as in the implementation of research artifacts. In addition, the deployment of the Disaster Portal during the 2007 southern California wildfires provided a wealth of information on what data is vital for ensuring effective communication between emergency response organizations and the public. In this one deployment, over 50,000 hits were made to the Disaster Portal site by the public and government agencies. In another successful example of living laboratory, the Peer-to-peer traffic information

dissemination testbed developed by UCSD, that now covers San Diego, Los Angeles, and Bay Area cities, has over 20,000 active users with an average thousand calls coming into our system every day.]

- 3 Develop integrative artifacts that will serve as a legacy for the RESCUE project, thus ensuring that the broader impacts of this five-year research program are realized.

[Progress in Yr 5: As a result of a re-focusing towards artifact development (which occurred in Year 3), the RESCUE project has produced 13 major artifacts with end-user partners. In addition to the artifacts, 31 research products were also developed. A number of the products are currently accessible through RESCUE websites, and are being used by both researchers and the general public (in the cases of the Peer-to-peer traffic Information System and Disaster Portal); a major effort will be made in the remainder of this year to ensure that more of these products and artifacts land in the hands of our government and industry partners.

- 4 Actively engage end-user community throughout the life of the project to validate the efficacy of the research and to serve as early adopters or testers of research products generated from RESCUE.

[Progress in Yr 5: The RESCUE project has worked closely with its Community Advisory Board, which is chaired by Ellis Stanley (the former Director of the Los Angeles Emergency Preparedness Department) to ensure that 1) end-users are aware of the RESCUE research program; 2) that each major project has a government partner to help identify research needs that will lead to implementation opportunities; 3) that RESCUE artifacts will have a home after completion of the current RESCUE project; 4) that proper tests and evaluations are conducted to ensure the efficacy of the research and its products; and 5) that partnerships developed during the current RESCUE project continue beyond Year 5. Some of the key organizations which have played a major role in this project are: the State of California; the County of Orange; the Cities of Los Angeles, San Diego, Irvine, Ontario and Champaign; the police, environment, health, and safety, and security departments at the University of California (Irvine and San Diego Campuses), and the U.S. Geological Survey. In addition, a number of industry partners have donated their time and equipment to the RESCUE project. As part of the continuity of Rescue and its artifacts, we have created a Center for Emergency Response Technologies (CERT) at UCI. The mission of CERT is to organize and direct research within the emergency response domain and to facilitate further technology adoption of Rescue technologies.]

- 5 Address the social, organizational, and cultural contexts in which technological solutions are adopted and implemented in order to better understand how appropriate technologies can be developed and transferred to users. Create awareness of issues in scientific and industrial communities through workshops, focus groups, panels and open testbeds.

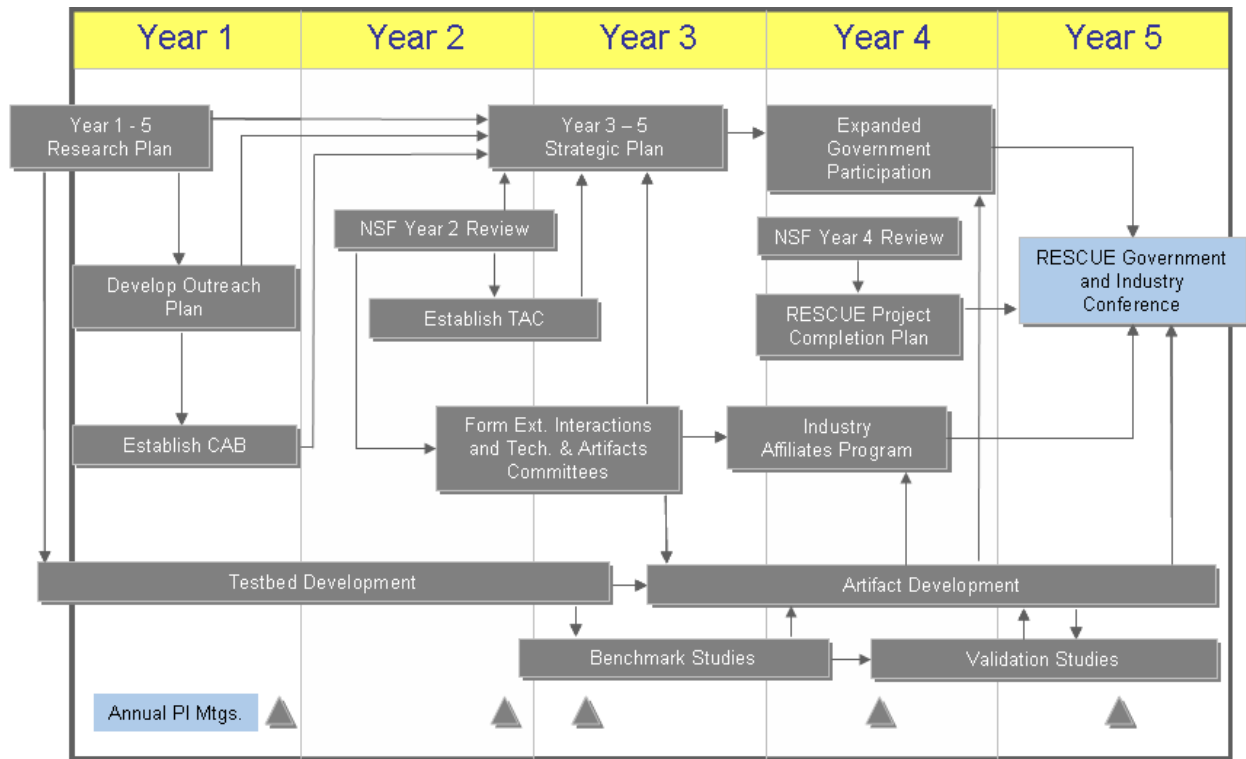
[Progress in Yr 5: The RESCUE program has had a core element that focuses on social, organizational, and cultural aspects of crisis response. This effort, led by Kathleen Tierney at the University of Colorado, has concentrated on the use of information and communication technologies among members of the public during disasters. Events that have been studied include the 2001 World Trade Center attacks, the 2004 Indian Ocean earthquake and tsunami, Hurricane Katrina, and the 2007 Virginia Tech shootings. Peer-to-peer communications behaviors in these events highlight the extent to which information and communication technologies are revolutionizing risk communication, information sharing, and collective sense-making with the public during extreme events.]

- 6 Actively engage a broad range of student populations through a multi-course interdisciplinary series on emergency response and focused research projects for graduate, undergraduate students. Create a concerted K-12 outreach effort through demos, lectures and internships. Leverage campus-level programs for underrepresented groups (e.g. California Alliance for Minority Participation and Women-in-CS) to actively recruit minority students.

[Progress in Yr 5: RESCUE continues to have a major impact on course curriculum throughout all the universities involved with the project. So far, 14 separate courses on crisis response topics have been developed and taught since the beginning of the RESCUE project. Some example classes include: System Artifacts geared towards First Responders; Special Topics in Information Technology for Homeland Security; and Issues in Crisis Response. In Year 5, the RESCUE Project had nine students graduate with either a PhD or Masters degree. In addition, RESCUE researchers have continued to reach out to the K-12 community by sponsoring high school interns and participating in campus events for high school students.]

1.2 Progress towards Project Milestones

Milestones where significant progress has been made to date are shown as grey boxes in Figure 1. Brief statements of this progress are also provided below.



The RESCUE Project has made significant progress in all key milestones for Year 5 with the exception of the Government and Industry Conference. Because of added emphasis in developing and testing project artifacts, and actively involving more end-users, the Executive Committee of the RESCUE Project decided to postpone the conference to after Year 5. Our current plan is to explore the possibility of holding this important conference towards the end of 2009. Should the project be successful in obtaining its one-year no-cost extension, this will be one of the final milestones of the current RESCUE Project.

As a result of the Strategic Plan developed in Year 3, the Executive Committee established a protocol requiring RESCUE artifacts to forge closer partnerships with the user community, specifically requesting advice on the design of artifacts, in addition to identifying opportunities for early adoption, testing and validation. Demonstrable progress has been achieved in the

majority of artifacts developed to date. Section 2.3 of the report describes this progress along with some of the key government partners that are involved with the RESCUE project.

Progress in each research project and artifact development effort is described in Sections 2.1 through 2.3 of this report. Key highlights include: 1) the emergence of cohesive multidisciplinary teams for each integrative research project; 2) the establishment of key partnerships with the user community for all the RESCUE artifacts; 3) the application of technologies in real-world environments, e.g., Disaster Portal in 2007 southern California wildfires and Peer-to-Peer traffic information dissemination testbed deployed in San Diego, Los Angeles/Orange Counties, and Bay Area; and 4) validation of research results using simulation methods, feedback from end-users, and application in actual events.

RESCUE is also continuing its push to engage industrial partnerships in the project. Numerous discussions, presentations, and demonstrations had been held with a number of companies including Boeing, Apani Network, IBM, Convera, Sun Microsystems, Printonix, Motorola, Ether2, IBM, Conexant, Yahoo, SBC, Asvaco, VDT, CDWG, Psion Teklogic and AMD, resulting in substantial equipment donations, new strategic alliances, and partnership discussions and/or non-disclosure agreements.

Efforts to increase participation of the user community in RESCUE include establishment of key partnerships with response organizations in the context of artifacts (e.g., City of Ontario for the Disaster Portal (www.DisasterPortal.org/Ontario) and expansion of the Community Advisory Board (e.g., addition of new members representing the Orange County Fire Authority & the Boeing Company). In consultation with Ellis Stanley, the Chair of the Community Advisory Board, the project has initiated a strategic expansion of CAB to increase the likelihood of national-level impact and project continuity beyond Year 5. Key researchers within RESCUE are in discussions with representatives from the Centers for Disease Control, the Federal Emergency Management Agency, the Office of Homeland Security for California, the Department of Homeland Security – Emergency Management, Los Angeles County – Department of Health Services/Bio-Terrorism/Preparedness and Response with regard to actively engaging their members in the project's CAB. RESCUE researchers have also deployed their technologies during several drills, including full-scale emergency response exercises.

Further efforts resulting in a strengthening of the industry/academic strategic partnerships include the creation of a school-based Center for Emergency Response Technologies (CERT). This Center, housed in the Donald Bren School of Information and Computer Sciences, was created as a vehicle for Rescue continuity and to facilitate our efforts to transfer technologies from the research laboratory and testbeds to the first responders and communities. Through CERT, we are pursuing licensing opportunities of Disaster portal technologies, most notably, with the city of Champaign Illinois and Rancho Cucamonga California.

2 RESEARCH PROGRESS IN YEAR 5

2.1 Research Projects

Project 1: Situational Awareness from Multimodal Input (SAMI)

Project Summary

The SAMI project is focused on research and technology development to realize the next-generation of situational awareness systems. Our ultimate goal is to develop an end-to-end situational awareness “engine” that can be used for particular situational awareness applications, primarily in the area of disaster response. Situational awareness, in the context of disaster response can be broadly viewed as consisting of the *past* (knowledge), *present* (comprehension), and *future* (prediction) state of the resources, infrastructure, entities (people), and the incident. Our research is aimed at addressing technical challenges in three key areas, namely information extraction and synthesis from raw sensory data, situational data management, and analysis and visualization technologies for decision making and support. The key technical approach we have investigated in extraction and synthesis is exploitation of (a) multimodality, and (b) semantic knowledge in extracting and interpreting data. The key challenges in situational data management addressed include representation and reasoning with uncertainty. In the context of data analysis our focus has been on understanding patterns of human behavior over time. Examples include analysis and understanding of Web access logs, event detection and prediction with vehicular traffic and accident data, and classifying human activities from low-cost observation modalities used for ubiquitous sensing such as RFID, video, et cetera.

Our work in the last several years has significantly increased our ability to extract and synthesize vital information from raw multi-modal data streams that are often available during disasters. The use of semantics has resulted in pioneering approaches that help to address traditionally complex data extraction problems. We have developed probabilistic models that learn patterns of human behavior that are hidden in time series of count data.

Many of the research outcomes of this project have been incorporated into RESCUE artifacts with Disaster Portal being the most prominent. For instance, technologies for entity disambiguation and extraction from text are the driving force behind the family reunification component of the portal. Likewise, predictive modeling techniques are the driving force behind occupancy forecasting.

Activities and Findings

As part of SAMI, solutions to many diverse and often traditionally difficult problems have been developed that are allowing research teams to realize their goal of achieving situational awareness using raw data. We have addressed many problems using different modalities including text, audio, sensor, and video and have applied these modalities to multi-modal *information extraction and synthesis*. In situational data management, we have developed event data models and data management systems at the level of events. We have also developed solutions to efficiently manage complex probabilistic data – motivated from the application of managing location uncertainty in situational awareness. The end use of situational awareness is through *analysis and visualization* in which we have created many advanced capabilities for end-user decision making. Some examples include a next-generation open Web scale GIS search engine motivated by the emergency response domain, a graph-

based semantic visualization framework for events, a predictive modeling framework that can provide route predictions in travel, a detector for anomalous events over sensor data, and advancements to real-time damage assessment technologies for disasters.

Information extraction and synthesis. In this area, the SAMI project examined what are traditionally considered complex problems in raw data extraction and synthesis. To address this problem, the research team adopted a novel approach where semantics were applied in a basic and fundamental way. The result was to significantly enhance our ability to more accurately extract higher quality information from text, audio, video, and sensor data. In text data extraction, we created a next-generation extraction framework that facilitates the integration of semantic information with traditionally-used text processing. In doing so, we took a broad approach in extracting information from text where we also addressed problems that result from incomplete or erroneous data or from the presence of redundant and ambiguous information. In SAMI, we have also pioneered the development of a relationship-based approach to address the data quality problem. The outcome is a graph-based disambiguation solution (a ReIDC relationship-based disambiguation framework) that is significantly more advanced when compared to other data cleansing approaches. The solution is scalable, domain-independent and uses a self-tuning framework that can be applied to several practical tasks such as the disambiguating entities on the Web – this task has direct benefit to the family reunification system of RESCUE (which is part of the Disaster Portal Artifact discussed later in the report). The theme of semantics is also present in audio information synthesis. Here, we have developed an approach that improves the quality of audio tags (provided in the context of images taken via a cell-phone) by using semantics in the form of frequencies and correlations across particular keywords. In each of the areas we have applied semantics to – text extraction, image tagging, entity disambiguation, location disambiguation, activity recognition – the fact that semantics can help interpret data is not surprising (it is rather expected). What is novel about our work and which we consider to be a breakthrough technology is that we have developed a principled probabilistic framework that is quite general (works in diverse application settings) and yet efficient (can be used for real-time or near real time interpretation of data in order to support situational awareness). The approach is based on representing domain knowledge in the form of a semantic entity relationship graph and mapping data interpretation/analysis tasks to an equivalent graph analysis task.

Our work in multi-modal event extraction has also resulted in an innovative approach to situational awareness in ubiquitous computing. The basic idea here is that in an environment with many audio communication devices, the location, timestamp, and user identification information are augmented along with the actual audio and shared in a reliable and intelligent fashion. Finally video event detection work has resulted in a comprehensive object and event tracking framework where one can track objects in video streams at a higher semantic object level (such as persons, buildings or vehicles of interest). This ability can be fundamental in developing robust situational awareness solutions.

Situational data management. Key outcomes in this area include earlier work on SAMI in which a semantic event model, E, was developed and where an event-based data management model was built on top of E. Furthermore, we developed an approach to interpret and manage location uncertainties (in a scalable manner) when the uncertainties arise because of colloquial references to location in speech or text. The work in this area overlaps with the SAT-Ware effort which is a comprehensive multi-modal sensor data collection and management platform. SAT-Ware is described in a separate section of the report.

Analysis and visualization. The focus here has been on pushing the state-of-the-art in analysis and visualization systems by developing powerful and intuitive-based solutions that specifically address improved decision making. Our earlier work resulted in a graph-based analysis framework to visualize and analyze relationships between events in particular disaster situations. An important component in this framework is GIS data. Our major contribution in this area is the development of a next-generation GIS search engine that scales to the open Web. Our work addresses some of the limitations encountered in locating relevant GIS data on the open Web by maintaining metadata about the various data sources. This work also includes histogram-based indexing techniques that allow the approach to easily scale to hundreds of thousands of GIS data sources on the open Web.

Our work on analysis techniques includes predictive modeling where sensor data (loop information) are used to predict route movements for people and cars and to determine whether certain travel activities represent anomalous or unusual movements. Similar techniques are used to predict occupancy levels of building exploiting people counter data. To address these questions, mixed Bayesian networks and high quality approximate sampling techniques are employed.

Specifically, in Year 5, we extended our work to link multiple people count sensors in a building. This model has shown to be effective at estimating the occupancy of the building; overcoming the bias introduced by sensor noise and missing and corrupted measurements which confound the simple counting techniques. Our results show that our proposed approach provides much more accurate estimates than the simpler baselines. A fault-tolerant extension of the original model was developed found to be effective at finding predictable traffic patterns and detecting unusual behavior when applied to a large network of over 1700 sensors and 100 million measurements. The model provided useful insights about an urban data set considered difficult to analyze. Both of these examples show how the original probabilistic framework can be extended and built upon without major changes to the implementation. This work has resulted in a general statistical framework for unsupervised learning of both recurrent patterns of human behavior over time as well as characteristics of unusual events (Ihler, Hutchins, Smyth, 2006, 2007; Ihler and Smyth, 2006).

Our research was evaluated using two real-world data sets – car counts from loop detectors located on California freeways, and people counts from optical sensors in a building on the campus of UCI. Validation sets were created with the help of the building events coordinator and from the baseball season schedule for the stadium near the loop detector. The efficacy of our models was largely judged by quantitatively assessing the accuracy of our model's predictions (e.g., checking how well the model predicted known events that were known for evaluation, but not known to the model during training) and simulating missing/noisy data and evaluating how well the model could recover such data.

While the algorithms and tools we have been developing have focused on analysis of occupancy within instrumented buildings and on loop sensor data, it can be used as a component of much larger situation assessment tools that can quickly and accurately assess how many people are in a given location from very noisy sensor measurements and also determine whether any unusual patterns are present or not. To this effect, we have initiated incorporation of our algorithms into the disaster portal artifact to provide first responders with occupancy information (useful in response planning) as well as to the public at large (useful for route planning in evacuations).

Finally, we have focused our research on developing methods to estimate damage and impacts in a real-time sense using sensor data from regional seismic networks. For the past decade, real-time information on seismic ground motions has been available through the U.S. Geological Survey. However, the ability to use this information to rapidly predict damage and impacts to communities has not been available. Using specially-developed techniques that allow the efficient application of GIS functions in a web-based environment, the research team has successfully developed a real-time earthquake damage assessment system for the entire State of California, which is currently being tested by several State and Federal agencies. This system – which has been integrated with the popular VirtualEarth interface – was used during the recent Chino Hills earthquake which occurred on July 26, 2008.

Products and Contributions

Much of the research under the SAMI project has resulted in software products that can be used, at least at a prototype level, by the research or end-user community. Some of these products have also matured to the level of being part of key RESCUE artifacts, such as the Disaster Portal. Key products include:

- **XAR** - XAR is an open-source text extraction system that allows integration of semantic information with existing off-the-shelf extractors. This is currently available with a Creative Commons license. An application of this extractor is a module for online internet monitoring which is part of the RESCUE Disaster Portal.
- **ReIDC** - ReIDC is a domain independent, self-tuning disambiguation system.
- **Multi-modal Audio Event Extraction Toolkit** - This toolkit can be used in ubiquitous environments where the sharing of auxiliary information such as location, times, and speaker identifications along with the audio information is needed.
- **Video Object and Event Tracking System** - This was developed by the UCSD CVVR group. It enables situational awareness by facilitating the tracking of semantic objects such as persons, buildings, vehicles, etc, in videos.
- **GAAL** - This product is a graphical analysis algebra toolkit that creates visualization relationships amongst events in an intuitive fashion.
- **The Travel Route Planner** - The route planner can predict travel routes and destinations for a person given his origin and other contextual information such as day of week, time of day, etc.
- **Anomalous Event Detection Toolkit** - A toolkit to detect unusual events based on an analysis of California Highway loop sensor data. An unusual traffic and highway events detector is being integrated into the Disaster Portal.
- **Prototype, Next-generation GIS Search Engine** - A search engine for GIS data that is significantly more comprehensive when compared to general purpose search engines or specialized one-stop portals. Work is currently underway to make this GIS search engine available as a module in the Disaster Portal.
- **InLET Real-time Disaster Damage Assessment Tool** - A rapid damage assessment tool that utilizes U.S. Geological Survey data on real-time strong ground motions for California earthquakes. Has been used in several recent earthquakes including the July 26, 2008 Chino Hills event.
- **Casualty Estimation Instrument** – A casualty estimation instrument that uses "scale-up" survey methods taken from social network research. The method uses reports from persons

tangentially affected or not directly affected by disasters to capture casualty information on third parties. The instrument can be employed in online surveys (analogous to the USGS Did You Feel It? system) for rapid assessment of casualties and other adverse outcomes (including injury, displacement, and property loss) in the immediate post-impact period.

- **Machine Learning Repository.** The loop detector data and people-count data used in our earlier RESCUE work (Ihler, Smyth, Hutchins, 2006, 2007) along with the validation data are also publicly available via the UCI machine learning repository at <http://archive.ics.uci.edu/ml/>.

Future Research Directions

Our work on SAMI has opened many new opportunities many of which will be pursued outside the scope of RESCUE. First and foremost, our work has established the possibility of a generic principled approach to exploiting semantics for online data interpretation tasks. While RESCUE has enabled us to conduct some preliminary work in this direction, achievement of the lofty goal requires significant new research. For instance, can one develop a systematic classification of diverse interpretation tasks which can then be addressed together, can spatial and temporal reasoning as well as reasoning in the form of rules be incorporated in the semantic-based approach to data interpretation, can one make such semantic approaches truly self-tuning, etc. Similar challenges can be identified with our work on statistical interpretation of sensor data.

While many of these questions that our research on SAMI has raised will be realized through other derivative projects, our focus in the near future will be two folds:

- (a) incremental steps towards extending our basic research. For instance, the next logical direction for statistical mining of sensor data would be to spatially link the traffic loop sensors. In the urban case study, we saw evidence of unusual event activities that had spatial signatures. This model would be used to determine the spatial and temporal signature of unusual activity, such as those caused by traffic accidents or other emergencies. Another interesting direction is performing large-scale dynamic population density estimation. This was one of our original goals for the traffic data set. The fault-tolerant extension of the model has helped us overcome some of the initial roadblocks we encountered with this important problem, especially in the context of situation assessment and emergency response. In the context of exploiting semantics for data interpretation, we would like to explore systematic ways to capturing and representing association and correlation amongst data. Exploiting ontological reasoning and incorporating that within our framework is also another obvious direction for future work.
- (b) focus on incorporation of the techniques we have developed so far into an end-to-end situational awareness system in the context of RESCUE artifact (viz. Disaster Portal). In addition, the research team will work closely with our government partners to ensure the effective use and adoption of the SAMI technology in real emergency response settings.

Project 2: Robust Networking and Information Collection

Project Summary

The main objective of this project is to provide research solutions that enable the restoration of computing, communication, and higher-layer services at a crisis site in a manner that focuses on the needs and opportunities that arise proximate to the crisis (in both time and space dimensions). Commercial systems are often based on assumptions that fall apart during a crisis when large-scale loss of power and destruction of antenna masts and servers are common. In addition, self-contained relief organizations that arrive at a crisis site often carry communication equipment that fail to interoperate, are inadequate for the needs at the scene, and may even interfere with each other making the task of forming an ad-hoc organization harder. In summary, the challenge is to compose a set of research solutions to assist in crisis response that is designed to serve the dynamically-evolving situation at the crisis site.

This project is organized into four sub-groups: 1) Theoretical Research; 2) Extreme Networks System (ENS); 3) Adaptive Information Collection System (AICS); and 4) System Integration. ENS, AICS and Rich Feeds (system integration) are the three artifacts developed by the project.

Activities and Findings

Theoretical Research Group. The main contributions during the past year include the development of new solutions for network design. Existing network optimization and throughput estimation are done in an asymptotic manner where the network size, or the number of nodes in the network, is expected to be infinitely high. However, in reality, the network size is finite and therefore, as part of RESCUE research, we proposed new methods for estimating the throughput and optimization of the network design for hierarchical wireless mesh networks. In addition, we also proposed packet-size-aware path setup mechanisms for wireless mesh networks. According to this scheme, different packets belonging to the same source destination pair can traverse different paths based on their sizes. We implemented our scheme as a Linux kernel module and experimentally showed the benefits in terms of the transmission delay and the throughput.

During the past year, we also proposed a new medium access control (MAC) protocol for Multiple Input Multiple Output (MIMO) ad hoc wireless networks. The key contribution is the use of a novel rank-based metric to obtain interference information, in order to enable multiple simultaneous transmissions and to make MAC decisions. Through analysis and simulation, we found that the proposed protocol significantly outperforms 802.11 MIMO and it obtained spatial degree of freedom utilization as high as 85%.

We modeled and studied supervisory network management (SNM) by developing an efficient communication model that can be used among response organizations. The SNM can ensure that only one update is sent at any given time and that it is received by all designated organizations before sending another one. In a limited resource network, this is to be done by a supervising entity that knows the resource availability. The system has been modeled from real observations made during the drills and from the UCSD emergency medical department, in order to improve human communication in crisis scenarios.

We created another important research result in collaboration with the CogNet project: a wireless network characterization strategy. In order to achieve efficient wireless network characterization, accurate traffic sampling strategies are required. Since the spectrum used for most wireless networks is divided into several channels, the existing count-based sampling methods are not scalable. Through experiments, we found that time-based sampling is an efficient and effective strategy for characterizing the wireless network environmental traffic. From our analysis using the Chi-Square test, we found that the Timer-driven, Time-based

sampling is more accurate than existing count-based sampling for both systematic and stratified sampling schemes.

We also focused localization solutions for two technological areas: first is localization using cellular network technology. For this, our research led to the proposal of a neural network-based algorithm that uses multi-path signal strengths to accurately locate a mobile user without a GPS receiver. Second, we developed a localization algorithm for ultra-wide band (UWB) technology. We proposed an artificial neural network-based localization algorithm to detect a single object in a sensor network and compared its performance to a Cramér-Rao bound and least squares estimator. We obtained good resolution when using our solution.

Extreme Networking System. Many research advances were made for several components and peripherals of the ENS. Notably, a new routing protocol for the CalMesh platform, called MACRT, was developed and successfully deployed. This replaced the prior spanning tree protocol. MACRT is a layer 2 (MAC) ad hoc on-demand routing protocol and was inspired by the popular layer 3 AODV protocol. But, as the name implies, it operates on layer 2 of the protocol stack, making the mesh nodes use MAC addresses to "route" within a mesh network. MACRT also incorporates several new functions, such as: (i) control message intercepting where it intercepts the 802.11 client management messages and uses these messages to help clients roam between Access Points (AP), (ii) the ETX (Expected Transmission Count) which is used as a link metric in the routing algorithm in order to achieve better throughput, (iii) a delay algorithm which introduces very short delays before "Route Requests" are forwarded, and (iv) a neighbor subsystem that maintains the connections to its adjacent nodes by using a bounded random walk model of the RSSI (Received Signal Strength Indication) values in order to filter out unstable neighbors.

Adaptive Information Collection System. As part of this project, we have created a fully automated peer-to-peer system in San Diego, Los Angeles, Orange County and the Bay Area (in northern California) that collects and disseminates highway incident information to the general public and to first responders. Commuters can report incidents via the system and other users who are calling in for traffic information will receive information about the incidents in their commute segment. We used the following strategies to validate and rank the user input: (i) access control mechanism which granted access only to frequent users (carefully fanned out over time) and (ii) a flagging mechanism using which other users can flag if they thought it was spam. More information on this system can be found in the description of Artifacts section of this report.

In addition to the traffic notification system, we enhanced the cellular phone-based location tracking system to develop the Participatory Symptom Network (PSNET). It is developed for real-time disease surveillance via mobile phones and it can be used: (i) by networks of health care providers and (ii) to monitor travelers as they cross geographical boundaries. In the first case, the application focuses on a patient monitoring system where doctors and nurses can keep track of sick patients, especially the elderly and infants. In the second case, monitoring of spreading diseases by visitors to and travelers in any country can be carried out.

System Integration Group. The system integration group mainly focuses on integrating various research products of the RESCUE project including Robust Networking and Information Collection. To meet these challenges, Rich Feeds' design is based on a Service Oriented Architecture (SOA) pattern called Rich Services, which delivers the benefits of SOA in a system-of-systems framework using an agile development process. This architecture allows the

addition of new data producers and consumers quickly and with low risk to existing functionality while providing clear paths to high scalability. More details can be found in the Artifacts section.

Products and Contributions

- **CalMesh2** - CalMesh2 is an improved version of our original CalMesh platform (<http://calmesh.calit2.net>) which is a mesh network platform that forms the basis for ENS.
- **CalNode** - CalNode (<http://calnode.calit2.net>) is a prototype Cognitive Network Access Point (CogNet AP), which has the unique capability to observe and learn from the network traffic in order to optimize itself. A CalNode can be deployed with no prior channel planning.
- **Rich Feeds** - Rich Feeds (<http://rescue.calit2.net>) is a system that demonstrates how unconventional data feeds and emergent data feeds can be captured, preserved, integrated, and exposed either in real time or after the fact. Rich Feeds promotes situational awareness during a disaster by integrating and displaying these feeds on a Google map in real time.
- **Web and Voice Portals for Peer-to-Peer Networking** - <http://traffic.calit2.net/sd>, <http://traffic.calit2.net/la>, and <http://traffic.calit2.net/bayarea>; (866) 500 0977 - San Diego, (888) 9Calit2 - Los Angeles and Orange Counties, and (888) 4Calit2 - Bay Area. This system is based upon the Calit2 Wireless Traffic system that relays customized highway incident information to the general public and to first responders.

Future Research Directions

In Year 6, our focus will be on refining, finalizing, improving, packaging, and disseminating the research solutions and artifacts that we created as part of this project.

Project 3: Policy-Driven Information Sharing (PISA)

Project Summary

The objective of PISA is to understand data sharing and privacy policies of organizations and individuals; and devise scalable IT solutions to represent and enforce such policies to enable seamless information sharing across all entities involved in a disaster. We are working to design, develop, and evaluate a flexible, customizable, dynamic, robust, scalable, policy-driven architecture for information sharing that ensures the right information flows to the right person at the right time with minimal manual human intervention and automated enforcement of information-sharing policies, all in the context of a particular disaster scenario: a derailment with chemical spill, fire, and threat of explosion in Champaign.

Activities and Findings

During this year, we addressed technical and sociological problems that were identified through the derailment crisis scenario that is the focus point for PISA. As discussed below, our main efforts during the past year were 1) continuing analysis of the sociology focus groups in Champaign; 2) a released version of the TrustBuilder2 software for runtime trust establishment; 3) work to provide user-friendly authorization facilities for crisis victims and their friends and family, and 4) work to provide information integration facilities that can amalgamate friends-and-family notices taken from grass-roots and government-sponsored family reunification web sites. We describe each of these in more detail below.

In August 2006, Rescue sociologists facilitated three focus groups for 28 first responders and other stakeholders in Champaign, exploring how the community's public safety and emergency management organizations would interact and communicate using technology in response to the derailment with chemical spill scenario. From the data collected, several key observations have been identified, including challenges and problems the community faces in this scenario, and technology solutions and suggestions. Analysis of the results of the focus groups is ongoing.

During the past year, we crawled three more autonomous sites and collected 46,000 more records. Currently the system has 76,000 records about missing people. One unique feature of the interface is that it can support interactive, fuzzy search. That is, the system can do search on the fly as the user types in more keywords. The system can find also records that may match user keywords approximately with minor differences. This feature is especially important since there are inconsistencies in crawled records from the Web sites, and the user may have limited knowledge about the missing people he is looking for. We are also adding an authorization mechanism to the system so that we can keep track of the users who submit information about missing people. In addition, we are adding a feature of displaying person information on a map using Google Maps.

Products and Contributions

- **Trustbuilder 2** - We have released the TrustBuilder2 runtime system for trust negotiation under a BSD-style open source license. Since its release, it has been downloaded over 170 times and is currently being used by several research laboratories. Last year, we designed, built, and evaluated an efficient solution to the trust negotiation policy compliance checker problem, called Clouseau, and incorporated it into TrustBuilder 2. Given some authorization policy p and a set C of credentials, Clouseau quickly determines all unique minimal subsets of C that satisfy p . In the last year, we have developed compilation procedures that translate policies written in the RT0, RT1, and WS-SecurityPolicy languages into a format suitable for analysis by Clouseau. We have also proven the correctness and completeness of these compilation procedures, which shows that Clouseau can safely and efficiently analyze policies written in existing language
- **User-friendly Authorization** - In response to confidentiality concerns identified in the derailment scenario for family and friends reunification, we worked to develop lightweight approaches for establishing trust across security domains. Victims need a way to ensure that messages they post are only read by the intended family members and friends, and vice versa. Obviously logins, passwords, PKI infrastructure, and other heavyweight authentication solutions are not practical in this context. SAW is our user-friendly alternative that eliminates passwords and their associated management headaches by leveraging popular messaging services, including email, text messages, pagers, and instant messaging. Additional server-side support integrates SAW with web technology (blogs, wikis, web servers) and browser toolbars for Firefox and Internet Explorer.
- **WARP** - During the past year, we concentrated our efforts on Wireless Authentication using Remote Passwords (WARP). Current single sign-on techniques, including our own Simple Authentication for the Web, require a user to directly contact a third party during authentication. These approaches are unsuitable for wireless access since the user does not have the network access necessary to contact a third party. WARP is a new in-band protocol that allows a user to prove to a wireless access point that he/she knows his/her password without the access point gaining access to his/her password or data that can be

used to launch an off-line attack on the password. WARP has the potential to be used beyond wireless access protocols, as well.

- **Data Integration Facilities** - Data from friends-and-family reunification web sites are extremely heterogeneous in terms of their structures, representations, file formats, and page layouts. A significant amount of effort is needed in order to bring the data into a structured database. Further, there are many missing values in the extracted data from these sites. These missing values make it harder to match queries to data. Due to the noisiness of the information, an integrated portal for friends-and-family web sites must support approximate query answering. At the start of this reporting year, we already had a working demo of our information integration technology that addresses these problems, available at <http://fr.ics.uci.edu/demo/>, with data from 16,000 missing person reports taken from three web sites.

Future Research Directions

Effective crisis response requires the ability to nimbly reconfigure the way that functions are accomplished. One aspect of this is the ability to quickly reassign powers and privileges, i.e., to change the authorizations in a system. It can be very hard to understand the short- and long-term ramifications of such changes. During the RESCUE project, we developed techniques to compile high-level authorization policies into a form that can be checked very efficiently at run time. We would like to leverage this approach to create a framework that enables one to quickly grasp the ramifications of changes in authorization policies. Such a policy analysis framework will have utility both during and after a crisis, by allowing post-mortem analysis to quickly zero in on actions that would not normally have been permitted. Such a framework would also be extremely helpful in any large organization during normal daily operations, as authorization policies are changed.

Project 4: Customized Dissemination in the Large

Project Summary

This project focuses on information that is disseminated to the public at large specifically to encourage self-protective actions, such as evacuation from dangerous areas, sheltering-in-place, and other actions designed to reduce exposure to natural and human-induced threats. Specifically, we have developed an understanding of the key factors in effective dissemination to the public and designed technology innovations to convey accurate and timely information to those who are actually at risk (or likely to be), while providing reassuring information to those who are not at risk and therefore do not need to take self-protective action. There are three key factors that pose significant challenges (social and technological) to effective information dissemination in crisis situations – variation in warning times, determining specificity of warning information to effectively communicate to different populations, and customization of the delivery process to reach the targeted populations in time over possibly failing infrastructures. Our approach to address these challenges is a focused multidisciplinary effort that 1) understands and utilizes the context in which the dissemination of information occurs to determine sources, recipients, channels of targeted messages and 2) develop technological solutions that can deliver appropriate and accessible information to the public rapidly. The ultimate objective is a set of next generation warning systems that can bring about an appropriate response, rather than an under- or over-response.

Activities and Findings

In the customized dissemination project, we are exploring three main research thrusts:

- 1) **Understanding Dissemination Context:** We have focused on information diffusion simulations to model the communication behavior of community members upon receipt of warning messages from our Crisis Alert system;
- 2) Supporting customization needs through flexible, timely, and scalable technologies; and
- 3) Scalable and Robust Delivery Infrastructure, both in wired and wireless networks

The key artifact of this project is the CrisisAlert System which allows fast and customized dissemination of alert messages through the Internet and incorporates some of the research findings achieved within this project.

During the past year, there has been significant progress in all three research fronts. Significant progress was also made in the implementation, testing and deployment of the CrisisAlert System in government organizations. Working with industry partner, Nokia, we have developed fast broadcasting services on Nokia devices equipped with multiple wireless interfaces. We are also in discussions with industry in the school-based communication domain (primarily Fonevia Inc.), through whom we expect to obtain participation of some Southern California school districts for the pilot study to be conducted in the coming year. We are also in discussion with several agencies to test and deploy versions of the CrisisAlert system.

Understanding Dissemination Context. In order to determine the effectiveness of the CrisisAlert system as well as the dissemination protocols that we have incorporated in it, our research team has performed a series of computer-based simulations. These simulations helped the research team assess the efficacy of the system, and in turn, how effective messaging to the population-at-large was. Using demographic data - as well as research from sociological studies of communication patterns and frequencies between various types of people - the simulations modeled the communication behavior of the community members upon receipt of warning messages from our Crisis Alert system. The simulation method was preferred over an initial real-world study for matters of practicality and speed, i.e., the information obtained about the information dissemination behaviors of simulated human actors was used to help inform the Crisis Alert system with far less time and money than would have been needed in a real-world study of these behaviors.

As the simulation model evolves, we can also factor in the geographic locations of the key study areas (e.g., schools) and establish various rules based on peoples' distances from important locations (e.g., schools, shelters, etc.) Overall, we are able to control and test for other effects in the information dissemination process that may be caused by factors such as sex, age, race, education level, and so forth. As we observe the various behaviors of the actors in the simulated world, we can further refine how the Crisis Alert system should be used.

Customized Information Dissemination in Large-scale Publish/Subscribe. The customized information dissemination can be provided through a scalable Publish/Subscribe (Pub/Sub) infrastructure. A hierarchical Pub/Sub system consisting of set of content brokers can provide scalability and fault tolerance for delivering relevant information to receivers. Existing approaches provide scalable Pub/Sub capabilities by forming efficient overlay networks to connect content brokers. Tree-based and DHT-based architectures are two of the most common architectures in forming broker overlay networks in distributed Pub/Sub systems. We have proposed a fast and robust broker overlay architecture that provides fault tolerance and speeds up the content dissemination process. The approach also balances the content

dissemination load among brokers which results in more robust and scalable infrastructure. Our approach achieves fault tolerance and scalability by organizing event brokers in clusters. Multiple inter-cluster links provide continuous availability of dissemination service in the presence of broker failure without requiring subscription retransmission or reconstruction of broker overlay. Furthermore, the proposed architecture balances broker load and provides a fast event dissemination infrastructure that significantly reduces subscription and publication dissemination traffic and load on event brokers. Our experimental results show that even in the presence of high failure rate in broker network, information dissemination is not interrupted and dissemination speed and load are not affected significantly.

Efficient management of subscriptions in Pub/Sub systems is also an important challenge. We have been working on a subscription subsumption approach to reduce the cost of subscription dissemination load among brokers. The proposed technique extends the subscription covering technique to prevent unnecessary dissemination of subscriptions that are subsumed by a set of previously stored subscriptions. We analyze the problem for d -dimensional content space and prove that the complexity of subscription subsumption evaluation is exponential. We then propose heuristics to exploit the subscription subsumption as much as possible and evaluate our proposed techniques through extensive simulations.

All of the above Pub/Sub systems ignore the heterogeneity in the receiver desired content format and deliver the published content to all the receivers in the same format. However, in real applications, users may desire to receive information in a variety of formats. For instance, users that receive information on handheld devices such as PDA or cell phone may desire to receive contents that are suitable for such devices. Therefore, a customized information dissemination system should not only deliver the information to relevant receivers but also it should customize the content format in such a way that it fits the receiver's context. Existing distributed Pub/Sub systems do not provide such service and disseminate information in the same format to all receivers. We have proposed a DHT-based Pub/Sub architecture that not only delivers content to the relevant users, but also customize the information for each user based on their desire. The customization is done through the content adaptation operators that accept content in one format and convert it into another format. Examples of content adaptation operators are content transcoders for multimedia information and content translation for multilingual receivers.

The important challenge in information customization in distributed Pub/Sub system is deciding about the place that the customization operations must be performed. Two straightforward options are to perform the required operations at the source or at the destination of the published information. However, such decisions may result in increased customization costs or increased network traffic or both. We define the customized Pub/Sub system as the problem of selecting the perfect location in the broker overlay for performing customization operators and showing that this problem is NP-Complete (Steiner Tree Problem). We then propose heuristics for selecting locations for content adaptation and compare the proposed techniques with the approaches that perform adaptation in the source or destination brokers. Our analysis shows that the proposed approaches reduce the computation cost resulting from customization operators along with the communication costs resulting from content transmission.

Scalable and Robust Delivery Infrastructure.

Scalable and Robust Information Dissemination in Wired Networks. During the past year, the research team has mainly focused on very short-term information dissemination, as is the case for seismic early warning where only few seconds are available before the earthquake strikes.

In this scenario, given that the information available is limited, the amount of data to disseminate to a large number of recipients is usually small. As in the case of Flash dissemination which we have already explored in previous years, information needs to be disseminated as fast as possible. Reliability is another key factor in this type of application; although the failure model that we have adopted does not include massive failures since the disaster has yet to strike yet and the network infrastructure is fairly stable.

We have studied existing protocols for group communication and have identified some that could apply to the early warning scenario. In particular, we considered various implementations of Application Layer Multicast and gossip protocols. We also identified some major issues associated with these protocols and have adapted some of them to our highly constrained scenario, where very low latency and high reliability are required. As part of our research in this project, we have developed a new protocol that merges both gossip and Application Layer Multicast advantages and exploits the knowledge of the group structure to minimize the overhead and the dissemination speed. We are currently testing it through simulations to show the advantages of this protocol in the early warning scenario as compared to others.

We are assuming a use-case scenario where schools throughout California are the possible recipients and groups of receivers can be built based on their locations. In fact, we have made the assumption that schools in the same geographic region (i.e. county or city) will need to receive the same early warning message and we have exploited this knowledge to speed up the dissemination process. Given this scenario, we can test the scalability of our protocol as the number of recipients increases.

In order to test our protocol in a complete and more realistic manner, we are considering a failure model that includes both independent random failures (pre-disaster) and geographically correlated failures that can happen as the disaster strikes. We are planning to determine to which extent to which our protocol can be used to effectively disseminate early warning under these failure assumptions. Finally, we will integrate this protocol in the Crisis Alert system, adding the capability to disseminate early warning or high priority messages.

Information Dissemination in Heterogeneous Wireless Environments. Wireless networks (e.g., cellular, Wi-Fi) extend the capabilities of landlines to warn and notify large numbers of people (both the public at large and first responders) in crisis situations. People with handheld devices (e.g., cell phones, PDAs) not only receive emergency alerts, but can also share warnings and other related information with each other via ad hoc networks. In this work, we study reliable, fast, and efficient dissemination of application-generated data in heterogeneous wireless networks.

Here, we address the reliable dissemination of application content data to mobile devices that may or may not have infrastructure support but rather are connected through ad hoc networks. This condition would apply to situations where mobile devices generate rich information (e.g., pictures, annotated maps, etc.) and this information must be disseminated to all other reachable devices nearby. We have developed an approach that decomposes the dissemination process into two subtasks: awareness assurance and data diffusion. We have developed and implemented the RADcast (Reliable Application Data Broadcast) protocol which provides reliable delivery of application data to all receivers regardless of whether network size knowledge is available. It is composed of an awareness diffusion sub-protocol, *Peddler*, and a data diffusion sub-protocol, *Pryer*. *Peddler* employs small messages called walkers which traverse the network to assure that the dissemination metadata reaches the recipients. *Pryer* provides guaranteed delivery of the content data to those dissemination-aware receivers

through concurrent push and pull. By intelligently integrating Peddler and Pryer, as well as smart exploitation of network size knowledge, RADcast supports three levels of reliability guarantee.

We have built a prototype system on Nokia N800 Internet Tablets, which forms Wi-Fi ad hoc networks among a group of devices in near proximity to each other. The system enables users to see each other's presence, do text messaging, file sharing and dissemination, video and audio streaming and talking. Communication can be one-to-one, one-to-many or one-to-all. The system has integrated our protocol, RADcast, for reliable content dissemination services.

Products and Contributions

- **Crisis Alert** - Crisis Alert is a software artifact for the dissemination of information to the general public during an emergency. The Crisis Alert system has been built with the goal of integrating research results directly into information dissemination technology. A guiding principle has been to respond to issues identified in the warning literature regarding over-response and under-response in crisis situations. To reach a greater part of the population and to overcome partial failure of the communication infrastructure, Crisis Alert delivers emergency notifications through different modalities, including CREW and the Early Warning protocol that we have developed.
- **Wireless Messaging Platforms** – This product is a prototype proximity messaging system built on Nokia N800 Internet Tablets, which forms Wi-Fi ad hoc networks among a group of devices in the vicinity. The system enables users to see each other's presence, do text messaging, file sharing and dissemination, video and audio streaming and talking. The system has been integrated into our protocol, RADcast, for reliable content dissemination services.

Future Research Directions

We are currently developing a prototype opportunistic messaging/dissemination system on Nokia S60 Series smart phones. The goal here is to enable opportunistic messaging and dissemination based on human encounters using short-range radio technologies such as Bluetooth and Wi-Fi on personal handheld devices. Other than for proof-of-concept purposes, this system will also serve as a platform for testing the effectiveness and efficiency of other opportunistic messaging techniques. Meanwhile it will enable collection of real-world data sets of human movement patterns and provide better understanding of the characteristics of human networks.

Project 5: Privacy Implications of Technology

Project Summary

Privacy concerns associated with the infusion of technology into real-world processes arise for a variety of reasons, including unexpected usage and/or misuse for purposes for which the technology was not originally intended. These concerns are further exacerbated by the natural ability of modern information technology to record and develop information about entities (individuals, organizations, groups) and their interactions with technologies – information that can be exploited in the future against the interests of those entities. Such concerns, if unaddressed, constitute barriers to technology adoption or worse, result in adopted technology being misused to the detriment of society. Our objective in the project has been to understand

privacy concerns in adopting technology from the social and cultural perspective, and to design socio-technological solutions to alleviate such concerns. We have focused on applications that are key to effective crisis management. For example, applications for situational awareness might involve personnel and resource tracking, data sharing between multiple individuals across several levels of hierarchy and authority, information integration across databases belonging to different organizations. While many of these applications have to integrate and work with existing systems and procedures across a variety of organizations, another ongoing effort is to build a “sentient” space from the ground up where privacy concerns are addressed right from the inception, trying to adhere to the principle of “minimal data collection.”

Activities and Findings

In this project, our goal has been to address the various privacy challenges in technology adoption for crisis response. We identified three main areas of focus: 1) observation/data collection systems, 2) data sharing systems, especially in the outsourcing model and 3) data dissemination systems. We made substantial headway in the design and implementation of observation and data sharing systems over the past year. We have also identified some important privacy-related issues that need to be addressed in data dissemination systems. We present a summary below of the work completed over the past year and the problems that have been identified for each of the three areas above for the coming year.

Observation & Data Collection Systems. In Year 5, we have made the following advances in observation and data collection systems related privacy research:

a). Observation systems refer to those systems that utilize a variety of sensing technologies to create situational awareness which is then utilized to provide new functionalities to individuals (as in pervasive computing applications) and/or monitor activities for security purposes (as in surveillance systems). In Year 4, we focused on designing a “privacy-preserving surveillance” application. The goal of such a system is to detect whenever an event of interest (from a pre-specified set of events) occurs in the pervasive space. The constraint being that it should not reveal the identity of individuals who were not involved in the events of interest. In fact, in the case of complex events (which unfold over an extended period of time), the identity of no individual is exposed until the final stage of the event occurs. We started working on a new privacy-aware pervasive space application called the RegionMonitor. In RegionMonitor users are allowed to pose queries that pertain to other individuals and regarding resources located within the pervasive space. The application allows individuals to specify their information needs as one-time or continuous queries and information sharing (about self) preferences as policies. For example, user A may pose a continuous query as follows, “Alert me whenever B and C are in the conference room before noon”. He may set the preference for this to be very high if the day is “Monday”. Similarly, A may specify the following rule amongst his information sharing policies, “Do not allow B (A’s boss) to know my location during the lunch hours unless it is an emergency but allow C to know my location at all times”. In RegionMonitor, policies are not only about privacy preferences of users, but about the generic utilization of the information about one’s information that is collected and made available to others in the system. To allow users to specify policies for information utilization, we designed a policy specification framework and a suitably expressive language. Specifically, it allows users to intuitively specify the utility of information they seek (provide) about (to) others. While most privacy/security implementations simply translate policies into access control rules, we in addition carry out inference analysis that considers all previously exposed information and a variety of background knowledge before making new information available to a user. Privacy violations encompass a security risk and thus, users express their willingness to take that risk by assigning negative utility values to such

information being exposed. We setup the problem as an optimization problem where the goal is to maximize the aggregate utility of the information released by the system. The resulting optimization problem turns out to be significantly complex, and its implementation is challenging especially in a real-time setting. We have developed a distributed, scalable solution technique for utility maximization in the RegionMonitor application. More description about the pervasive infrastructure, SATware on which RegionMonitor is built, is provided in the artifacts section of this report.

b). Data generated in dynamic environments like pervasive spaces and location-based applications is critical for understanding the complex manner in which human movement and information access patterns intermingle. Understanding the dynamics of such complex systems is essential for designing better and more efficient applications in the future. But making the data generated in such environments available for research raises a plethora of privacy concerns regarding the storage, processing and sharing of the information contained therein. In Year 4, we developed a probabilistic framework for analyzing privacy threats in dynamic location-based applications. In Year 5, we extended these concepts to the problem of privacy-preserving trajectory data release. We defined a novel notion of location-privacy for trajectory data that takes into consideration adversarial knowledge like background information and domain knowledge. The goal as we mention above, is to release aggregate level trajectory information while preventing the adversary from inferring in the specific locations that individuals occupied under a variety of attack scenarios. In this work, we also identified generic utility measures that cover broad classes of queries and developed location-generalization algorithms using trajectory clustering schemes. The generalization algorithm is an optimization procedure that tries to maximize the utility of the released data while simultaneously ensuring that all the privacy constraints are met.

c). Our extensive work on dynamic systems has significantly enhanced our understanding of the nature of privacy in such systems. We have been able to identify common issues that occur across a variety of dynamic applications that need to address privacy concerns. These range from the very definitions of privacy, the types of privacy constraints, the utility measure and finally to the solution techniques that are employed. While we have made some progress towards developing a framework to reason about and deal with privacy in dynamic systems, much work still remains in completing the formalization. This will be a focus in the coming year.

Data Outsourcing for Easy Sharing & Mobile Access. In Years three and four, we developed two systems (DataGuard and DataVault) which provided users with secure remote storage and remote network drives, respectively. The main benefits of such systems were ease of mobile access and sharing. In Year 5, we continued to work on making the two systems more stable and adding new features. Significant amongst them is the search functionality that was built. The goal was to support multidimensional search over the encrypted data stored on the server-side. A dimension here refers to one of the many attributes that the data might have and a user is able to search over the set of encrypted objects by specifying selection criteria for one or more of these attributes (e.g., keywords for text attributes, range of numeric values or dates etc.). While previous work on search over encrypted data has assumed the availability of a compliant server (i.e., a server that implements the proposed algorithms/protocols), we explore a more challenging problem, that of supporting secure search utilizing only the vanilla keyword search functionalities provided by web-services today. Our solution is a middleware that carries out the data transformation, indexing, cryptographic operations, key-management, query translation etc. that are required for providing such functionality.

Products and Contributions

- **RegionMonitor** - RegionMonitor is a SATWare application that allows users to pose queries pertaining to other individuals and resources in the pervasive space. Each user is able to specify his information-usage policies which the system attempts to optimize. In RegionMonitor, policies are not only about privacy preferences of users, but about the generic utilization of the information about the individual that is collected and made available to other users of the system.
- **DataGuard**: (<http://dataguard.ics.uci.edu>) DataGuard middleware allows users to outsource their file system to heterogeneous data storage providers available on the Internet. DataGuard also preserves data confidentiality and integrity of outsourced information using cryptographic techniques. DataGuard is fundamentally designed to be easy to use to the average user.
- **DataVault**: (<http://www.ics.uci.edu/~rjammala/DataVault>) DataVault is an architecture designed for Web users that allows them to securely access their data from any machine connected to the Internet and also lets them selectively share their data with trusted peers. The DataVault architecture is built on the outsourced database model (ODB), where clients/users outsource their database to a remote service provider who delivers data management services such backup, recovery, transportability and data sharing. In DataVault, the service provider is untrusted. The confidentiality and integrity of the user's data is preserved using cryptographic techniques. The service provider manages encrypted data. DataVault utilizes a novel PKI infrastructure and encrypted storage model that allow data sharing to take place via an untrusted server.

Future Research Directions

While we have made significant contributions to two of the three subprojects mentioned above, we have not made significant progress in addressing the privacy requirements in data dissemination systems. In the coming year, we intend to focus our efforts significantly in this direction. One interesting problem in this domain is that of privacy-preserving content dissemination in a publish/subscribe network. The problem is to efficiently match information requests (queries) from users to data being produced in the network without revealing too much information about the information seeker.

Besides the data dissemination work, there are several challenges remaining in the other two subprojects. To name a few, we have issues of properly analyzing inference risks in observation systems and strategies to deal with future inference. For any piece of information that the observation system exposes to a user, there exists, potentially, some auxiliary information that the user could use to infer some private piece of information. Properly analyzing the inference risk of exposing some information to a user is still an open challenge in observation systems. Moreover, in an observation system, inference avoidance is also more challenging than in other systems due to the fact that users can change their queries at any time. That is, based on what the observation system answers, users can change their continuous queries. Thus, the observation system needs to take into account such possibility in order to maximize the utility of the information being released and minimize the privacy violation risks.

In the data sharing work, there are several issues that need to be addressed. For instance, more flexible sharing policies could be implemented than in the current version of DataGuard. Search optimization is another area which we are currently working on. A few of the issues we are currently working on are compression of search indices, search on compressed indices, selection of suitable encryption functions that can handle various data formats etc.

Project 6: MetaSIM

Project Summary

MetaSIM is a web-based collection of simulation tools developed to test the efficacy of new and emerging information technologies within the context of natural and manmade disasters, where the level of effectiveness can be determined for each technology developed. MetaSIM currently incorporates three simulators: 1) Crisis simulator InLET; 2) Transportation simulator, and 3) Simulator for agent based modeling (Drillsim).

Activities and Findings

In Year 5, a website has been developed in HTML and stored in the backend database to produce web pages on-the-fly through Java script. The web pages call the various simulators and allow users to define parameters for the various simulations. The parameters are saved in user-specified scenarios and the simulations are run through the interface. After each run, the results are stored in the database and the website calls and displays intermediate and final results. The application is supported by a DB2 database (with Spatial Extender) and an ArcIMS server that stores geographic data and creates all the different GIS layers through a standard interface. For the first version of MetaSIM, the geographic extent of the data for implementation was Los Angeles and the Orange County area. The transportation simulator and Drillsim have been implemented for a much more focused area centered around the University of California, Irvine campus.

Several model refinements were implemented for individual simulators. For the Crisis Simulator/ InLET, user-defined parameters to run a custom scenario were included as a part of the meta-simulation. The definition of the evacuation scenario for DrillSim was also implemented. Multi-floor including indoor-outdoor agent evacuation was completed for DrillSim. For the transportation simulator, time synchronization and data exchange with pedestrian network using a Whiteboard database was completed. The protocol to inform the MetaSIM testbed along with technology assumptions was explored for Adaptive Cellular Networking System.

A new Relational Spatial Data Model was developed to overcome the challenges associated with varied spatial data and multiple simulator integration within MetaSIM. This new standard for model integration enables MetaSIM as a testbed for technology testing by addressing the integration of multiple geographies and simulators, and a variety of spatial data models (vector, raster, or network).

Current online mapping applications such as Virtual Earth and Google Earth offer a rich representation of information layers including base layers of road, aerial and satellite imagery. Technologies for data access, sharing and distribution securely over the internet make it possible to push information to a large population at a very rapid rate. All of these factors combined with the reduction in hardware costs have created an environment where an online loss estimation program like InLET provides greater flexibility to the disaster management and response community. Implemented over the popular online Virtual Earth mapping interface, InLET results are overlaid on a rich layer of Virtual Earth data and imagery.

MetaSIM is also a collection of plug-and-play simulation tools connected by a database. In its final form, the definition of inputs, outputs, timing, and scale, and the results of each simulation component will be available for iterative use by each of the other simulation models. Registering and synchronizing transactions between various simulation engines and assuring

proper use of scale will be addressed by the data exchange architecture and the time synchronization module.

MetaSIM is developed with open software architecture to enable modules to share data in real time. The platform and protocol designed for MetaSIM's data exchange support modular and extensible integration of simulators for the scientific, engineering, and emergency response communities.

Products and Contributions

- **Crisis Simulator/ InLET** -The Crisis Simulator currently simulates an earthquake event and estimates damage and casualties for Los Angeles and Orange counties. The crisis simulator integrates the earthquake loss estimation components of InLET; the Internet based Loss Estimation Tool.
- **DrillSim** - DrillSim is an agent-based activity simulator that models human behavior at the individual or micro level. DrillSim tests IT solutions by modeling situation awareness and provides it to the agent to react accordingly. For example, an early warning system might be used to modify the timing of agent evacuation. Micro-level activity modeling provides the ability to mimic agent behavior in crisis, as well as interactions between people during crisis, thereby providing a more robust framework for integrating responses to information and technology.
- **Transportation Simulator** – The transportation simulator consists of an integrated model of simplified quasi-dynamic traffic assignments, and a destination choice model. Information that becomes available through IT solutions is simulated through parameters, such as subscription to routing support information via cell phone or email, information arrival time and update frequency, system credibility and acceptance, to reduce uncertainties associated with decision making when evacuating a congested network.
- **GIS Applet for Visualization** - A GIS applet has been developed for the crisis simulator for visualization of the different geographic data layers and the simulation results.
- **InLET version 2** - A more refined crisis simulator, InLET v2, has been developed for integration into the MetaSIM framework to replace the existing damage and loss modeling simulator. InLET has been expanded from its previous implementation for Los Angeles-Orange County to cover earthquake hazards throughout California. InLET provides rapid results by simplifying the FEMA/HAZUS damage functions and porting the algorithms to a database environment. Given the advances in online mapping and the availability of GIS data through various online platforms, the visualization component of InLET has been implemented over the popular online Virtual Earth interface.

Future Research Directions

Various options are being evaluated for deploying the core components of MetaSIM individually as well as a complete system. We seek to leverage the capabilities of METASIM and are working to transform it into a tool that can be used for multiple purposes. Some possible applications include:

- 1 Loss estimation and decision support tool for public agencies;
- 2 Integration into Web Portals such as the City of Ontario's portal www.disaster.org;
- 3 Training tool for first responders;
- 4 Modeling tool for online services; and

5 Online disaster game for the gaming industry

2.2 Other RESCUE-Related Research

Privacy Concerns in Instrumented Spaces. This research study investigated privacy concerns in large-scale, highly instrumented environments. The primary goal of the study was to identify individuals' privacy concerns as they are immersed in Rescue-related research and sensing technologies. This two-phase research project consisted of: 1) observing any changes in pro-social behavior in a closed-loop feedback environment consisting of RFID, optical sensors, acoustic sensors, motes, and a "feedback portal;" and 2) a social sciences survey specifically addressing the privacy concerns in such a highly instrumented environment.

Business Continuity and Disaster Recovery Leadership. This research project (a dissertation study performed by a RESCUE team member) investigated correlations between the depth and breadth of business continuity/disaster recovery planning and that of leadership frame orientations and leadership characteristics. The RESCUE social sciences team assisted in creating the web-based instrument that was sent to Southern California academic institutions. The resulting data from the instrument is part of the Rescue Data Repository.

2.3 RESCUE Artifacts

Extreme Networking System (ENS): The primary goal of the ENS system is to develop an efficient, reliable, and scalable network infrastructure to aid and support emergency response activities. Significant technical advances were made for several components and peripherals of the ENS system. Notably, a new routing protocol called MACRT was developed for the CalMesh platform; it was successfully introduced and tested during the Winter/Spring 2008. In addition, the CalMesh hardware was overhauled to a new version with a more capable Linux platform and faster WiFi cards (we now run a 500 MHz CPU and use 802.11g rates on Atheros cards). We call this updated platform CalMesh2 and it can be either hosted in the same rugged aluminum cases as the first version, or in a smaller form factor enclosure that is also designed for outdoor use.

The routing protocol, denoted MACRT as in MAC Routing, was first developed for the ENS system as a cross-layer routing protocol in order to solve the drawbacks of the prior spanning tree protocol. The MACRT was used to test and verify functionality of the protocol. After the prototyping phase, the per-packet intense operation of the protocol was moved into the Linux Kernel space in order to achieve the desired performance in terms of throughput and latency. MACRT is a layer 2 (MAC) ad hoc on-demand routing protocol and was inspired by the popular layer 3 (IP) AODV protocol. However, as the name suggests, it operates on layer 2 of the OSI stack, which means that the mesh nodes use MAC addresses to "route" between source and destination mesh nodes. Similar to AODV, every new packet flow is preceded by a short "Route Request" message flooded in the mesh network, which will invoke a "Route Reply" from the source node once it is received. The nodes along the "best" path will form the route for the new packet flow in the following way: the flooded "Route Request" packets all set up a reverse route in all the nodes it passes on its way from the source to the destination node. The destination sends the "Route Reply" message back along the route from which the "best Route Request" was received (it carries a link cost aggregated along the route).

One of the primary successes this year was the tight integration among the multiple components of the ENS system, which was showcased and tested in full-scale drills (as Rich

Feeds on a non-uniform tiled display [NUTSO]). The integration provided a more cohesive, interoperating infrastructure. This was successfully demonstrated when we showcased nearly a dozen of our research prototypes, tools and devices during the two emergency drills in which we participated this year: UCSD Campus Drill (full-scale exercise with an active-shooter scenario), October 16, 2007 and an MMST full-scale drill, dubbed Operation Silver Bullet (in South Bay, San Diego, CA, January 24, 2008 (unique scenario with dual incidents situation, the first time for us, and the local MMST).

Peer-to-Peer Adaptive Information Collection System: This artifact is a fully automated peer-to-peer system (<http://traffic.calit2.net>) in San Diego, Los Angeles, and the Bay Area (in northern California) that collects and relays highway incident information to the general public and to first responders. Though government agencies and the private sector have some of the basic data needed for effective highway incident collection, the means to effectively disseminate the data in an intelligent manner (i.e., delivery of relevant and timely information to the right segment) is lacking. Typically the data is disseminated in a broadcast mode, with unacceptable latencies. We have used San Diego, Los Angeles, Orange County and the Bay Area as a test bed to develop, deploy and test the above-mentioned system to empower the general public (in particular, commuters) to act as human sensors and relay information about incidents ranging from wild fires, mudslides and other major accidents to the general public and to the 911 control center. The system can be accessed simply by making a phone call and is based on speech recognition. We have learned from past experience that the general public will not adopt such a system if you inject a new phone number during the time of a disaster (such as the San Diego wild fires last fall). The system should be available consistently and regularly, disseminating information that is valuable to the public on an everyday basis. The system currently provides personalized real-time traffic information to the commuters via cell phone. We have modified this system so that commuters can report incidents 24x7, including the time, location, severity and the urgency of the event. We will analyze the data for validity and populate the events in a GIS database. Other commuters calling in, hear these events if they happen to fall in their commute segment. An example of a more severe case would be the San Diego wild fires spreading to I-15 resulting in a shutdown of the freeway. If one reports such an event, due to the severity of the event, the system will trigger an alert to all the users, to avoid that region of the freeway. As of today, this system handles more than 20,000 active users with every day a 1000 calls coming in.

Rich Feeds System: Rich Feeds (<http://rescue.calit2.net/>) is a system that demonstrates how unconventional data feeds and emergent data feeds can be captured, preserved, integrated, and exposed either in real time or after the fact. Rich Feeds promotes situational awareness during a disaster by integrating and displaying these feeds on a Google map in real time. One of the primary uses of this system is for the integration of various research products of the RESCUE project including Robust Networking and Information Collection. It demonstrated how unconventional data feeds, including research feeds from Robust Networking and Information Collection project and Situational Awareness from the Multi-model Input project, can be integrated and emergent data feeds can be captured, preserved, integrated, and exposed in either real time or after the fact. To meet these challenges, Rich Feeds' design is based on a Service Oriented Architecture (SOA) pattern called Rich Services, which delivers the benefits of SOA in a system-of-systems framework using an agile development process. Rich Feeds' service-oriented architecture allows for the addition of new data producers and consumers quickly and with low risk to existing functionality, while providing clear paths to high scalability. As a result, Rich Feeds provides users with the opportunity to integrate research and emergent feeds to create novel presentations and gain novel insights both in emergency and research settings.

Disaster Portal: The Disaster Portal (www.disasterportal.org) is an easily customizable web portal and set of component applications which can be used by first-responders to provide the public with real-time access to information related to disasters and emergency situations in their community. Current features include a situation overview with interactive maps, announcements and press notifications, emergency shelter status, and tools for family reunification and donation management. The Disaster Portal dramatically improves communication between first-responders/government agencies and the public, allowing for rapid dissemination of information to a wide audience.

The development of the Disaster Portal is based on two primary considerations. While we aim to provide practical applications and services of immediate utility to citizens and emergency managers, we also strive to significantly leverage many relevant pieces of IT research within RESCUE. The advanced technologies that are currently incorporated into the Disaster Portal include components for customizable alerting, family reunification, scalable load handling, unusual event detection and internet information monitoring.

Major achievements made by the Disaster Portal project during Year 5 include the demonstration of several new research elements within the application modules, deployment of the system at the City of Ontario, California, and packaging of the system for further deployments with additional cities. Infrastructure scalability was an area of active development during the past year. During disasters, information portals such as the Disaster Portal commonly experience flash loads due to very large numbers of users attempting to access the portal simultaneously. Two approaches were incorporated in the Disaster Portal to address this issue. The RESCUE Flashback web server creates and utilizes a peer-to-peer infrastructure and transmits much of the flash load back to end-user browsers which are dynamically recruited to share some of the server load. The second approach developed utilizes cloud-computing services such as Amazon's Elastic Computing Cloud. The capability allows for dynamic provisioning of virtual machine instances hosting the Disaster Portal system in order to load balance flash traffic.

Also during the past year, the PISA Simple Authentication for the Web (SAW) technology was integrated into the Disaster Portal making it possible for users to create accounts without having to remember a password for the portal.

SAW: The PISA project's SAW authentication technology eliminates the need for users to remember a password. SAW accomplishes this by leveraging popular messaging services, including email, text messages, pagers, and instant messaging. This technology was implemented as a deployment option within the Disaster Portal application.

TrustBuilder2: To address the systems issues associated with the adoption of trust negotiation as a flexible authorization approach for virtual organizations, we have designed, implemented, evaluated, and released the TrustBuilder2 framework for trust negotiation. TrustBuilder2 is a flexible framework designed to allow researchers to quickly prototype and experiment with various approaches to the trust negotiation process. In TrustBuilder2, the primary components of a trust negotiation system are represented using abstract interfaces, any of which can be overridden or extended by users of the system. TrustBuilder2 is also agnostic with respect to the formats of credentials and policies used during the negotiation process; support for new credential formats and policy languages can easily be added by extending the appropriate classes in TrustBuilder2's abstract type hierarchy. The flexibility of TrustBuilder2, along with its support for the interposition of user-defined plug-ins at communication points between system

components, not only enables users to rapidly implement support for new features, but also provides a framework within which the trade-offs between various system configurations can be quantitatively analyzed. The full source code of the TrustBuilder2 framework is available under a BSD-style open source license at <http://dais.cs.uiuc.edu/tn>, and has been downloaded more than 200 times, and is currently being used by several research laboratories.

Within the TrustBuilder2 framework, we have investigated the design of efficient policy compliance checkers. Given a policy p and a set C of credentials, a compliance checker is responsible for determining one or more minimal subsets of C that satisfy p . A compliance checker capable of finding all such satisfying sets of credentials would afford its users significant strategic and privacy-preservation benefits. Such a compliance checker would give users a complete view of the next-step state space for the negotiation, thereby allowing them to choose a locally-optimal satisfying set to disclose (e.g., the least sensitive satisfying set). Because many trust management policy languages are based on declarative logics (e.g., Datalog), traditional top-down theorem-proving approaches have long been used to solve the compliance checker problem. These types of theorem provers find at most one proof for a particular conclusion; as a result, compliance checkers built using this approach find at most one satisfying set of credentials at a time. Unfortunately, results in the research literature show that the time overheads of using these types of compliance checkers to find all satisfying sets of credentials grow exponentially in the size of the union of the satisfying sets discovered. This approach is not only asymptotically undesirable, but also impractical, as it can take seconds to find only a small number of satisfying sets. In an effort to increase the efficiency of trust negotiation systems, as well as enable participants to make more informed decisions, we developed the Clouseau compliance checker.

Clouseau: Clouseau takes a non-traditional approach to the type of theorem proving used in policy compliance checking. Specifically, Clouseau compiles trust negotiation authorization policies into patterns, translates credentials into abstracted objects, and then leverages efficient pattern matching algorithms developed by the artificial intelligence community to determine all satisfying sets of credentials for a particular policy. The running time of this algorithm scales as $O(NA)$, where N is the number of satisfying sets and A is the average size of each satisfying set; this is much more efficient than existing approaches. In more practical terms, our implementation of Clouseau - which runs inside of the TrustBuilder2 framework - can find hundreds of satisfying sets per second and can find all satisfying sets of credentials for practical policies in about the time required for one hard disk access. However, the model theories used to prove the correctness of more traditional compliance checkers do not map directly onto this pattern matching approach in the general case. To account for this, we have proven the correctness of Clouseau under a broad class of circumstances. Specifically, we have described compilation procedures for translating trust management policies written in the RT0, RT1, and WS-SecurityPolicy policy languages into a format suitable for analysis by Clouseau. We then proved that when operating on these compiled policies, Clouseau finds all satisfying sets for the original policies, finds no extraneous satisfying sets, and does all of this much faster than existing approaches.

SATware system: SATware (<http://ics.uci.edu/~projects/SATware>) is a multimodal sensor data stream querying, analysis, and transformation middleware that aims at realizing a sensor-based observation system that was originally designed as a RESCUE artifact to serve as a testbed for research on privacy and on situational awareness. As part of RESCUE, we have instrumented about a 1/3 of the campus including the Calit2 building as well as the newly constructed Bren Hall with a variety of sensors including cameras, audio sensors, motes with temperature sensors, mobile devices with GPS technologies, radios, people counters, etc. This sensing

infrastructure (that we refer to as Responsphere) is used to implement and monitor a variety of campus drills and experiments that test and validate RESCUE technologies.

In contrast with previous work on distributed stream systems and distributed sensor databases, in SATware our architectural objective is to hide the raw data being sensed and instead provide a semantically meaningful view of a geographic region and its entities to the applications. To this end, we modeled sentient spaces with two main layers of abstraction: the physical layer that consists of sensors, actuators, networks, and computing and storage nodes, and the semantic layer which captures the semantics of the real-world being sensed. Such a semantic level is modeled using data modeling techniques such as ER diagrams. Sensor data is translated through the concept of “virtual sensors” into meaningful events that are modeled as changes in attributes associated with the entities and relationships.

Based on the above semantic approach, we designed and implemented SATware which now serves as a middleware architecture connecting over 300 diverse types of sensors (including over 200 cameras) spread all over UCI. In addition, SATware runs on mobile platforms and provides interfaces to associate new sensors, new operators, as well as a powerful declarative query language based on SQL for writing applications. SATware supports a policy language for encoding privacy policies which are specified at the semantic level and translated to the sensor level. Furthermore, SATware supports SATmonitor and SATscheduler components that monitor and automatically schedule resources to optimize quality of data capture. In addition, techniques have been incorporated to exploit semantics to dynamically calibrate/recalibrate sensors.

Progress in SATware has witnessed an advance of SATware from being an artifact for demonstration of privacy & situational awareness research to being an important derivative research goal for RESCUE in itself. There are currently two PhD students working on SATware and the SATware team has applied for additional extramural funds to sustain the ongoing research in the area. In particular, SATware has transitioned into becoming the software “glue” for the recently funded DHS FEMA project at UCI.

Crisis Alert: an Artifact for Customized Dissemination: The Crisis Alert system artifact has been built during the past year with the goal of integrating research directly into information dissemination technology and to respond to the issues identified in the warning literature regarding over-response and under-response in crisis situations. In fact, Crisis Alert has the ability to send emergency notifications that are customized to the needs of each recipient. Crisis Alert also contains rich information such as maps of the area, location of the open shelters closer to the recipient's location, current state of hospitals and their address and contact information and they can be automatically created by the system according to a set of rules defined during the risk-knowledge phase of deployment of a warning system.

To reach a greater part of the population and to overcome the partial failure of communication infrastructure, Crisis Alert delivers emergency notifications through different modalities. In addition, Crisis Alert takes advantage of the emergency response plan of each organization, integrating social networks in the emergency dissemination process.

The main goal of this year for the Crisis Alert system has been to test it in real scenarios and incorporate the findings due to these tests. The Ontario drill has provided useful inputs to improve the usability of the system, regarding the case where no policies have been identified for the current emergency. In order to face this situation, the policy language has been enriched with the concept of “protective action,” allowing the emergency personnel to specify policies that can be applied to different and unpredictable events. The drill has also highlighted the need for

defining the concept of group of events. In fact, when a major disaster strikes, it usually generates a set of emergencies that are related to the major one but that can be of different nature and require different countermeasures. In this case, the population response could be improved if complete information is provided by the authorities through a single notification or update.

We are also trying to validate the Crisis Alert system through a series of pilot studies that involve schools and educational institutions. There are multiple purposes for these studies: to deploy the software prototype that we already have into a test scenario, to gather feedback from both the emergency personnel and the alert recipients involved, and to compare the information learned from this feedback to information obtained from actual drills. In this study, we would like to take advantage of the infrastructure that has already been put into place by Fonevia, which is used on a day-to-day base for disseminating information from schools to parents during non-disaster times. In fact, if people are already familiar with the dissemination system, it is more likely to obtain a prompt reaction when a warning is issued. Furthermore, we will supplement our technology testing with a simulation framework that will help us understand the alert dissemination in the whole community. Given knowledge of the geographies, policies and protocols – we can conduct a what-if analysis of the speed at which information can be spread in the community given different technology and usage scenarios.

Finally, in the next months we are planning a workshop with representatives from Southern California schools and school districts to collect information about existing warning systems, processes and procedures for emergency warnings and alerts to schools. This information will help us design the next generation of technologies and processes to help educational institutions better prepare for disasters and effectively respond in real-time to emergencies.

MetaSim: METASIM, both a project and an artifact, is a web-based collection of simulation tools developed to test the efficacy of new and emerging information technologies within the context of natural and manmade disasters, where the level of effectiveness can be determined for each technology developed. METASIM currently incorporates three simulators: 1) Crisis simulator InLET, 2) Transportation simulator, and 3) simulator for agent based modeling (Drillsim). For a detailed discussion of this artifact, see Section 2.1.

FlashBack P2P Web Server: Flashback 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. 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.

Responsphere: Responsphere is an IT infrastructure set of testbeds that incorporates a multidisciplinary approach to emergency response drawing from academia, government, and private enterprise. The IT infrastructure is used to test the efficacy of RESCUE technologies and extract meaningful metrics regarding those technologies. We view the testbeds as proving grounds for disruptive technology. During Year 5, the Responsphere infrastructure was significantly extended through the UCI campus. Several RESCUE artifacts (e.g., Disaster Portal and MetaSim) were tested and refined throughout the various Responsphere test-beds. Additionally, several first response activities such as drills and technology deployments were either conducted within Responsphere or facilitated by Responsphere.

Data Repository: RESCUE has created a publicly available, with proper IRB clearance, repository of disaster response related data sets (<http://rescue-ibm.calit2.uci.edu/datasets>).

These data sets are an extensive collection of information including (but not limited to): GIS maps, 911 calls, news broadcasts, mote telemetry data, and data gathered from emergency response exercises.

3 BROADER IMPACT

3.1 Community Outreach

Outreach to First Responders. Partnerships and collaboration with first responders have been an important strategy in RESCUE from the very beginning of the project. During this reporting year, the Disaster Portal was utilized by the City of Ontario during the 2007 California wildfires. During the fires, the system received over 50,000 hits and was used to post more than 20 announcements with maps, emergency sheltering status information, and to notify media organizations of news updates. The success of the Disaster Portal at the City of Ontario has led to interest from a number of additional cities, and discussions are ongoing to coordinate additional deployments. These include the Cities of Rancho Cucamonga, Aliso Viejo, Rancho Santa Margarita, and Orange; all in Southern California, and also the City of Champaign, Illinois.

Another example of successful implementation is the use of InLET (The Internet-based Loss Estimation Tool) which is built on MetaSIM. InLET currently uses ground motion information from the U.S. Geological Survey's real-time ShakeCast System. This system produces a regional "ShakeMap" that is based on a network of hundreds of ground-based sensors that record ground motions for earthquakes of moderate and large magnitudes (see <http://earthquake.usgs.gov/resources/software/shakecast/>). The RESCUE team has been working closely with the U.S. Geological Survey to develop a loss estimation system that takes the ShakeMap data and automatically imports that information into InLET after every earthquake. InLET then uses that information to estimate buildings losses and number of casualties in near real-time. InLET was used after the recent Chino Hills earthquake that affected the southern California region on July 26, 2008. The research team is currently in discussions with the California Governor's Office of Emergency Services (OES) and the Federal Emergency Management Agency to deploy InLET statewide in California.

Other RESCUE artifacts are continually being tested and validated during drills several of our participating universities, including UC Irvine, UC San Diego and the University of Illinois at Urbana-Champaign. For example, the EH&S division at UCI, the UCI campus police, and the Orange County Fire Authority (OCFA) continue to use RESCUE technologies at variety of emergency response drills. Included in these drills during Year 5 were two "active shooter" campus exercises (one at UCSD and one at UCI) that featured RESCUE dissemination technologies as well as networking technologies. Other drills included several evacuation exercises, including a very large-scale exercise involving the San Diego Metropolitan Medical Strike Team (MMST) and the UC San Diego Police and Emergency Services departments on the UCSD campus.

Industry Interactions. We continue to actively cultivate industrial affiliates as well as government partners as part of our outreach mission. RESCUE is working closely with developers at D-Link Inc. to refine a next generation of optical and acoustic sensors. We are also working with Caltrans to take their real-time data feeds of traffic counts and use this information as input into a number of our projects. Finally, the RESCUE project has been quite successful at embedding our graduates into companies such as Nokia, Jive Software Systems, VM Ware, and Motorola.

Broader Community: The RESCUE Disaster Portal and the Peer-to-Peer traffic system are examples of successful outreach to the broader user community. Citizens are empowered (through RESCUE Technology) to not only make day-to-day decisions regarding their commutes but make potential life saving decisions by virtue of the information disseminated to them. During the California wildfires, the Disaster Portal received over 5,000 separate “hits” from citizens seeking situational information.

Several high-level meetings have been conducted with State and Federal officials in emergency management to discuss opportunities for deployment of RESCUE technologies. For example, RESCUE team members have demonstrated the capabilities of MetaSIM to the GIS and Earthquake/Tsunami program managers at the California State Office of Emergency Services and the Mitigation Directorate of DHS/FEMA.

3.2 Education Outreach

Graduate and Undergraduate Education. RESCUE continues to have an impact on course curriculum throughout all the universities involved in the project. Specific courses and class projects have been designed to have a direct tie to the research being done at RESCUE. Details of UCI’s ongoing RESCUE seminar series can be found at <http://www.itr-rescue.org/outreach/seminars.php>. Throughout the year, the RESCUE project has encouraged undergraduate students to be a part of ongoing research through individual study courses, honors courses, the NSF-funded California Alliance for Minority Program (CAMP), Calit2’s SURF-IT program, and undergraduate research appointments.

Rescue students who have graduated in the current program year include: Ping Zhou, PhD; Yiming Ma, PhD; Dawit Seid, PhD; Mayur Deshpande, PhD; Iosif Lazaridis, PhD; Qi Han, PhD; Christopher Davison, PhD; Jonathan Cristoforeti, MS; and Allesandro Gighi, MS.

At UCI, we also participated in the Calit2-sponsored SURF-IT program for undergraduate research. The program is a 10-week summer program and RESCUE hosted Sean Li. His research focused on “Automatic Inference of Anomalous Events from (California) Traffic Patterns. In addition to the SURF-IT program we taught the following Rescue related courses at UCI: ICS 192, ICS 214A, ICS 214B, ICS 215, ICS 203A, ICS 278, ICS 199, ICS 290, ICS 280, ICS 299.

At UCSD, Ramesh Rao, BS Manoj, Rajesh Hegde, Javier Rodriguez, Don Kimball and Per Johansson supervised more than a dozen and a half undergraduate and graduate group design courses in electrical or mechanical engineering and computer science during the 2007-2008 academic year. Ramesh Rao advised an undergraduate student for three quarters in the academic year 2007-2008 for a self study course which significantly overlaps with RESCUE research area. Rajesh Hegde and BS Manoj taught graduate/undergraduate level courses where topics directly from RESCUE research were covered. Ingolf Krüger covered the RESCUE integration architecture in two computer science & engineering graduate courses he taught.

K-12 Education. RESCUE continues to reach out to the K-12 community by sponsoring high school interns and participating in campus events for high school students. During the 2007-2008 academic year, UCSD hosted 12 high school senior interns from the Preuss School, a charter school under the San Diego Unified School District whose mission is to provide an

intensive college preparatory curriculum to low-income student populations and to improve educational practices in grades 6-12.

In the upcoming months, we are planning a workshop with representatives from Southern California schools and school districts to collect information about existing warning systems, processes and procedures for emergency warnings and alerts to schools. This information will help us design the next generation of technologies and processes to help educational institutions better prepare for disasters and effectively respond in real-time to emergencies.

Internships and Student Exchange Programs. Students participated in the creation and evaluation of several RESCUE technologies, several students from the senior group design courses in electrical or mechanical engineering went on to intern on various subprojects. At UCSD, Prof. Ramesh Rao hosted several visiting graduate students, including one from the University of Karlsruhe, Germany, who is studying disaster response there.

Academic Community: Rescue researchers and technologists from UCSD campus gave a number of keynote addresses and invited talks. These addresses provide the Responsphere team the opportunity to engage a number of stakeholders (Government, industry, Academia, and First Responders) within the emergency response domain. Researchers at Rescue have hosted or participated in a number of events, sessions, and workshops throughout Year 5. For detailed information regarding these events, please reference the “latest news” section under www.itr-rescue.org.

Of particular outreach interest within the academic and broader community was the launching of a school-based center at the Donald Bren School of Information and Computer Sciences (UCI). The center, The Center for Emergency Response Technologies (CERT – www.cert.ics.uci.edu), seeks to coordinate research efforts within the emergency response domain. Additionally, the center is hosting monthly workshops on various important topics within emergency response (e.g., medical perspectives on emergency response technology).

4 FUTURE PLANS

Specific research that will be conducted in the current year is presented in Section 2.1. In general, research activities will focus on completing current research tasks, as presented in the RESCUE Strategic Plan (developed in Year 3), and in working with government partners to implement RESCUE artifacts within their respective agencies. In the coming year we plan to (a) finalize our research artifacts, devices, systems and technologies, (b) leverage new opportunities created by recent disasters, and (c) complete outreach and deployment efforts that were initiated in the Strategic Plan.

Extending and Finalizing Artifacts: Although the development of research artifacts was not a part of our original proposal, they became an important and successful way for us to verify and test our research results through first-responder drills and tests with our research partners (both industrial and government/public safety). This has required more interaction and feedback from government partners to demonstrate the efficacy of research results. We plan to focus our efforts on engineering, research, or dissemination for finalizing and packaging the research artifacts. Packaging the research artifacts is to do so in a way that will enable other users, whether they are within RESCUE project institutions, or project collaborators and even groups outside the project, to use the artifacts and adapt them to their own applications. To achieve this, we will finalize the research artifacts that we have developed, by creating a methodology to select the ones we want to preserve, completing documentation on those, and placing the artifacts within a public computing environment or virtual computing environment (such as the Amazon EC2 cloud and storage) so they can be easily accessed and used by others. The following examples indicate specific research integration tasks that will be undertaken in the context of RESCUE artifacts include:

1. *Metasim*: This artifact is an integration architecture for simulating response activities in combination with technology simulation and crisis simulation. Using Metasim, one can simulate an earthquake, its impact, as well as a response activity such as evacuation as well as simulation of technologies such as cellular communication, notification, etc. to study how incorporation of technologies could help in improving response. We currently have 2 students pursuing their PhD research on this area -- Vidhya Balasubramanian is exploring issues related to scalability of MetaSIM and Leila Jalaji is exploring an integration architecture (motivated by existing software architectures based on computational reflection). Both are expected to make very significant progress into their PhDs and we seek an extension to both support their work as well as to incorporate the research into the existing Metasim artifact.

2. *Domain Semantics for Situational Awareness*: Over the past 2 years we have developed a completely new approach to exploiting "domain semantics" in the form of integrity constraints and implicit semantics hidden in data for the purpose of interpreting raw sensory information. We have applied such a semantic approach for variety of tasks including data cleaning, speech understanding, and text extraction. Two PhD students working on their dissertations in this area include: Zheng Zhang and Pouria Pirazadeh both of whom are expected to make significant progress over the coming year. Similar to Metasim progress, we are expecting the students to culminate their research and incorporate the research ideas into existing artifacts -- in this particular case, into disaster portal. With the help of the semantic approach, the capability of disaster portal will be significantly enhanced since now it will be able to interpret RSS feeds of other situational data -- a capability that our first responder partners have identified as quite significant.

3. *Disaster Portal Customization:* The Disaster Portal has emerged as amongst the most vital legacies of RESCUE. There are ongoing dialogs with four different cities in Orange County as well as with the City of Champaign for licensing and usage of the Portal technologies. Each of these installations requires a certain degree of customization to the system to fit the specific needs of each organization. In the coming year, we intend to (a) make the artifact more robust, (b) self manageable -- today, if any organization needs to use disaster portal, our team/programmers are required to install and customize the artifact for the organization, (c) support additional features which our first responder partners have indicated would be a significant value added. In addition to the semantics-based extraction work described above, we will incorporate the predictive modeling research our team had conducted over the past 5 years related to predicting occupancy and unexpected events based on spatio-temporal data (flow of people/vehicles using loop sensor data) into disaster portal. Also, it will also allow us to integrate the Web based loss estimation tool we have developed (called INLET) into the portal technology to provide a comprehensive solution.

Leveraging New Opportunities: Adapting research activities to incorporate the observations and findings of several recent disasters, e.g., the 2007 Southern California Wildfires, the 2008 Sichuan, and the China earthquake. New data are being used to calibrate or validate specific research hypotheses. These areas of research include calibrating population density and flow estimation, efficacy of a mediated information dissemination, understanding the role of the internet in informing citizens, as well as the role of speech technologies for situational awareness and for customized alerts. Successful development of some research activities have led to unexpected opportunities to leverage RESCUE artifacts in real-world crises. The Disaster Portal has been used in the 2004 Indian Ocean Earthquake and Tsunami, and the 2007 Wildfires. Likewise, the traffic system developed at San Diego has been expanded to cover other Southern California regions.

Completing Outreach Efforts: Several of our key government partners have been busy responding to several large California disasters, including the southern California wildfires at the end of 2007 and the wildfires in northern and central California which just ended. Therefore, our ability to interact with them to complete the validation phase of some projects has been limited. For instance, one of the key planned activities was the organization of a series of stakeholder workshops for early warning and alerts with various agencies, representatives from the school district at various levels, and industry. While we have had successful meetings earlier this year, the follow-on meeting which will lay out a city-wide testbed for validating the alert system technology has been delayed due to unavailability of key participants. The meeting has finally been planned for sometime in September 2008 with the goal of designing a validation study for the experiment that will be conducted either towards the end of this year or very early next year. Barring no new disasters, we hope to complete this phase of research (validation) by Spring of next year.

APPENDIX A: PUBLICATIONS IN YEAR 5

1. B.S,Manoj; Zhou,Ping; Rao,Ramesh. On Adding Link Dimensional Dynamism to CSMA/CA based MAC protocols, Proceedings of IEEE Globecom 2007,2007-11
2. BS,Manoj; Zorzi,Michele; Rao,Ramesh R.. On the Use of Higher Layer Information for Cognitive Networking, IEEE Globecom 2007,2007-11
3. B.S.,Manoj; Kimball,Don; Rao,Ramesh. On the Viability of Wireless Mesh Networks as a Next Generation Wireless Networking Alternative, Proceedings of ICAST 2007,2007-12
4. Hellewell,Phillip L.; van der Horst,Timothy W.; Seamons,Kent E.. Extensible Pre-Authentication in Kerberos, 23rd Annual Computer Security Applications Conference (ACSAC),2007-12
5. Hutchins,J.; Ihler,A.; Smyth,P. Modeling Count Data from Multiple Sensors: A Building Occupancy Model, Computational Advances in Multisensor Adaptive Processing (CAMSAP) 2007,2007-12
6. Ihler,A; Hutchins,J; Smyth,P. Learning to detect events with Markov-modulated Poisson processes, ACM Transactions on Knowledge Discovery from Data, Volume 1 , Issue 3,2007-12
7. Shivappa,Shankar T.; Rao,Bhaskar D.; Trivedi,Mohan M.. An Iterative Decoding Algorithm for Fusion of Multimodal Information, EURASIP Journal on Advances in Signal Processing,2008
8. Lee,Adam J.; Winslett,Marianne. Enforcing Safety and Consistency Constraints in Policy-Based Authorization Systems, ACM Transactions on Information and System Security (TISSEC),2008
9. Mitra,Soumyadeb; Winslett,Marianne; Hsu,Windsor W.. Query-based Partitioning of Documents and Indexes for Information Lifecycle Management, ACM SIGMOD Conference on Management of Data,2008
10. Hariharan,Ramaswamy; Hore,Bijit; Mehrotra,Sharad. Discovering GIS Sources on the Web using Summaries, 2008 IEEE Joint Conference on Digital Libraries,2008
11. Jafarpour,Hojjat; Hore,Bijit; Mehrotra,Sharad; Venkatasubramanian,Nalini. Subscription Subsumption Evaluation for Content-based Publish/Subscribe Systems, none,2008
12. Jin,Yuzhe; Rao,Bhaskar D... Performance Limits of Matching Pursuit Algorithms, 2008 IEEE International Symposium on Information Theory,2008
13. Jin,Yuzhe; Rao,Bhaskar D.. Insights into the stable recovery of sparse solutions in the overcomplete representations using network information theory, 2008 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2008
14. Lazaridis,Iosif; Mehrotra,Sharad. Aggregate Queries, Progressi ve Approximate, Encyclopedia of GIS,2008: 11-16
15. Ma,Yiming; Kalashnikov,Dmitri V.; Mehrotra,Sharad. Towards Managing Uncertain Spatial Information for Situational Awareness Applications, IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), 2008

16. Hegde,Rajesh M.; Kurinawan,Joseph; Rao,Bhaskar D.. On the Design and Prototype Implementation of A Multi modal Situation Aware System, IEEE trans. on Multimedia (In Review),2008
17. Gogate,Vibhav; Dechter,Rina. Approximate Solution Sampling (and Counting) on AND/OR search spaces, 14th International Conference on Principles and Practice of Constraint Programming (CP), 2008
18. Gogate,Vibhav; Dechter,Rina. AND/OR Importance Sampling, In 24th Conference on Uncertainty in Artificial Intelligence (UAI), 2008
19. Gogate,Vibhav; Dechter,Rina. Studies in Solution Sampling, In 23rd Conference on Artificial Intelligence (AAAI),2008
20. Zhou,Ping; B.S.,Manoj; Rao,Ramesh. On Optimizing Non-Asymptotic Throughput of Wireless Mesh Networks, Proceedings of IEEE CCNC 2008, 2008-01
21. B.S.,Manoj; Nuggehalli,Pavan; Rao,Ramesh. On the Use of Information Sharing in Wireless Networks, Proceedings of IEEE CCNC 2008, 2008-01
22. Vaisenberg,Ronen; Ji,Shengyue; Hore,Bijit; Mehrotra,Sharad; Venkatasubramaniam,Nalini. Exploiting Semantics for Sensor Re-Calibration in Event Detection Systems, Multimedia Computing and Networking 2008, 2008-01
23. Zhou,Ping. On optimizing wireless mesh networks, Thesis, University of California, San Diego, 2008-01
24. Tamma,Bheemarjuna R.; B. S.,Manoj; Rao,Ramesh. On the Accuracy of Sampling Schemes for Wireless Network Characterization, Proceedings of IEEE WCNC 2008, 2008-03
25. B.S.,Manoj; Zhou,Ping; Rao,Ramesh. Dynamic Adaptation of CSMA/CA MAC protocol for Wide Area Wireless Mesh Networks, Elsevier Computer Communications, 2008-03
26. Lee,Adam J.; Winslett,Marianne. Towards an Efficient and Language-Agnostic Compliance Checker for Trust Negotiation Systems, 2008 ACM Symposium on Information, Communcation, and Computer Security (ASIACCS), 2008-03
27. Lee,Adam J.; Winslett,Marianne. Towards Standards-Compliant Trust Negotiation for Web Services (Extended Version), UIUC Dept. of Computer Science Technical Report, 2008-03
28. Mitra,Soumyadeb; Winslett,Marianne; Borisov,Nikita. Deleting Index Entries from Compliance Storage, EDBT, 2008-03
29. Harding,Andrew; van der Horst,Timothy W.; Seamons,Kent E.. Wireless Access using Remote Passwords, 1st ACM Conference on Wireless Network Security (WiSec), 2008-03
30. Jammalamadaka,Ravi Chandra; Mehrotra,Sharad; Seamons,Kent E.; Venkatasubramaniam,Nalini. Middleware Support for Protecting Personal Data from Web based Data Services, none, 2008-03
31. Jammalamadaka,Ravi Chandra; Mehrotra,Sharad; Seamons,Kent E.; Venkatasubramaniam,Nalini. Middleware Support for Protecting Personal Data from Web based Data Services, none, 2008-03

32. Jammalamadaka,Ravi Chandra; Gamboni,Roberto; Mehrotra,Sharad; Seamons,Kent E.; Venkatasubramanian,Nalini. iDataGuard: Middleward Providing a Secure Network Drive Interface to Untrusted Internet Data Storage, 2008-03
33. Jammalamadaka,Ravi Chandra; Gamboni,Roberto; Mehrotra,Sharad; Seamons,Kent E.; Venkatasubramanian,Nalini. iDataGuard: Middleward Providing a Secure Network Drive Interface to Untrusted Internet Data Storage, 2008-03
34. Bonsi,Valentina; Hore,Bijit; Mehrotra,Sharad; Yu,Xingbo. Clustering Techniques for Trajectory Anonymization, 2008-03
35. Kim,Minyoung; Massaguer,Daniel; Dutt,Nikil; Mehrotra,Sharad; Ren,Shangping; Stehr,Mark-Oliver; Talcott,Carolyn; Venkatasubramanian,Nalini. A Semantic Framework for Reconfiguration of Instrumented Cyber Physical Spaces, Second Workshop for Event-based Semantics 2008 (WEBS 2008) (held in conjunction with IEEE RTAS 2008 and as part of the Cyber-Physical Systems Week (CPSWEEK)), 2008-04
36. Shivappa,Shankar T.; Rao,Bhaskar D.; Trivedi,Mohan M.. Multimodal Information Fusion Using the Iterative Decoding Algorithm and Its Application to Audio-Visual Speech Recognition, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '08), 2008-04
37. Kazuhiro Minami, Adam Lee, Marianne Winslett, and Nikita Borisov. Secure Aggregation in a Publish-subscribe system, UIUC Technical Report UIUCDCS-R-2008-2968, May, 2008.
38. Lickfett,Jay; Ashish,Naveen; Mehrotra,Sharad; Venkatasubramanian,Nalini; Green,Jacob. The RESCUE Disaster Portal for Disasters and Emergency Response, Proceedings of the 5th International ISCRAM Conference, Washington DC, USA, 2008-05
39. Dilmaghani,Raheleh B.; Rao,Ramesh R.. A Wireless Mesh Infrastructure Deployment with Application for Emergency Scenarios, International conference on Information Systems for Crisis Response and Management (ISCRAM) 2008 Conference, Washington D.C., May 2008, 2008-05
40. Poster-IEEE Homeland Security, Waltham, Boston, MA, May, 2008.
41. Demchak,Barry; Krueger,Ingolf. Rich Feeds for RESCUE, Proceedings of the 5th International ISCRAM Conference, 2008-05
42. Lee,Adam J.; Winslett,Marianne. Towards Standards-Compliant Trust Negotiation for Web Services, Proceedings of the Joint iTrust and PST Conferences on Privacy, Trust Management, and Security (IFIPTM 2008), 2008-06
43. Sruthi Bandhakavi, William H. Winsborough, Marianne Winslett. A Trust Management Approach for Flexible Policy Management in Security-Typed Languages. 21st IEEE Computer Security Foundations Symposium, Jun, 2008
44. Walker,Daniel D.; Mercer,Eric G.; Seamons,Kent E.. Or Best Offer: A Privacy Policy Negotiation Protocol, IEEE Workshop on Policies for Distributed Systems and Networks (Policy 2008), 2008-06
45. Ergüt,Salih; Rao,Ramesh R.; Dural,Özgür. Localization via Multipath Strengths in a CDMA2000 Cellular Network Using Neural Networks, World Conference on Computational Intelligence 2008, 2008-06-01

46. Kalashnikov,Dmitri V.; Nuray-Turan,Rabia; Mehrotra,Sharad. Towards Breaking the Quality Curse. A Web-Querying Approach to Web People Search., the Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, 2008-07
47. Jafarpour,Hojjat; Mehrotra,Sharad; Venkatasubramanian,Nalini. A Fast and Robust Content-based Publish/Subscribe Architecture, IEEE Network Computing and Applications, 2008-07
48. Hutchins,Jon; Ihler,Alexander; Smyth,Padhraic. Probabilistic analysis of a large-scale urban traffic data set, Second International Workshop on Knowledge Discovery from Sensor Data (ACM SIGKDD Conference, KDD-08), 2008-08
49. Bhorkar,Abhijeet; BS,Manoj. Selection Diversity based MAC Protocol for MIMO Ad hoc Wireless Networks, IEEE Globecom 2008, 2008-11
50. Jafarpour,Hojjat; Hore,Bijit; Mehrotra,Sharad; Venkatasubramanian,Nalini. Subscription Subsumption Evaluation for Content-based Publish/Subscribe Systems, ACM/IFIP/USENIX Middleware 2008, 2008-12
51. Kalashnikov,Dmitri V.; Chen,Zhaoqi; Mehrotra,Sharad; Nuray-Turan,Rabia. Web People Search via Connection Analysis, 2008-01-01
52. Kalashnikov,Dmitri V.; Chen,Zhaoqi; Mehrotra,Sharad; Nuray-Turan,Rabia. Web People Search via Connection Analysis, The IEEE Transactions on Knowledge and Data Engineering,2008-01-01
53. Krueger,Ingolf; Meisinger,Michael; Menarini,Massimiliano; Pasco,Stephen. Under Review.
54. Butts, Carter T. `A Relational Event Framework for Social Action. Sociological Methodology, forthcoming. 2008
55. Butts, Carter T. Social Networks: A Methodological Introduction. Asian Journal of Social Psychology, 11(1), 13-41, 2008
56. Butts, Carter T. Network Structure In Spatial Contexts. Department of Anthropology Colloquium, Stanford University. Stanford, CA. 2007-10.
57. Butts, Carter T. Graph Theoretic Convexity. 28th Sunbelt Network Conference (INSNA), St. Petersburg, FL. 2008-01
58. Butts, Carter T. Latent Focus Models for Network Structure. Fourth Joint Japan-North America Mathematical Sociology Conference, Redondo Beach, CA. 2008-05.
59. Butts, Carter T. On Having Something In Common: Latent Two-Mode Models for Social Networks. Networks in Political Science Conference, Cambridge, MA. 2008-06
60. Butts, Carter T. Parameterizing Exponential Family Models for Random Graphs: Current Methods and New Directions." Invited presentation, ``Formation and Function of Social Networks" Mini-Conference, 2008 SIAM Meeting, San Diego, CA. 2008-07
61. Butts, Carter T. and Acton, Ryan M. Spatial Analysis of Social Networks. ASA Meeting, Boston, MA. 2008-09
62. Lind, Benjamin E.; Tirado, Miguel; Butts, Carter T.; and Petrescu-Prahova, Miruna. Brokerage Roles in Disaster Response: Organizational Mediation in the Wake of

- Hurricane Katrina. International Journal of Emergency Management. 2008. forthcoming
63. Petrescu-Prahova, Miruna and Butts, Carter T. Emergent Coordinators in the World Trade Center Disaster." International Journal of Mass Emergencies and Disasters. 2008. forthcoming
 64. Eguchi, Ronald T. and Hope A. Seligson, Loss Estimation Models and Metrics, Chapter 6 in Risk Assessment, Modeling and Decision Support, Editors: A. Bostrom, S. French and S. Gottlieb. 2008
 65. Eguchi, R.T., Huyck, C.K., Ghosh, S. and Adams, B.J. The Application of Remote Sensing Technologies for Disaster Management, Proceedings of the 14th World Conference on Earthquake Engineering, Oct 12-17th, Beijing China. 2008
 66. Cutter, S.L., Emrich, C.T., Adams, B.J., Huyck, C.K. and Eguchi, R.T., New information technologies in management," Chapter 14 in Emergency Management, Principles and Practice for Local Government, Second Edition, Editors, Waugh, W.L. and Tierney, K. 2007
 67. Sarabandi, P., Kiremidjian, A.S. and Eguchi R.T., Application of Statistical Inferencing Techniques in Building Inventory Compilation, Proceedings of the 5th International Workshop on Remote Sensing Applications to Natural Hazards, Sept. 10-11th, Washington D.C. 2007
 68. Friedland, C.J., Adams, B.J. and Levitan, M.L., Results of Neighborhood Level Analysis of Structural Storm Surge Damage to Residential Structures, Proceedings of the 5th International Workshop on Remote Sensing Applications to Natural Hazards, Sept. 10-11th, Washington D.C. 2007
 69. McMillan, A., Adams, B.J., Morley, J.G., and Chesworth, S, A tiered reconnaissance approach towards flood monitoring, utilizing multi-source radar and optical imagery, Proceedings of the 5th International Workshop on Remote Sensing Applications to Natural Hazards, Sept. 10-11th, Washington D.C. 2007
 70. Rathje, E.M. and Adams, B.J.,, An Evaluation of Earthquake Damage Scales used in Remote Sensing Studies of the 2003 Bam, Iran Earthquake, Proceedings of the 5th International Workshop on Remote Sensing Applications to Natural Hazards, Sept. 10-11th, Washington D.C. 2007
 71. Friedland, C., Adams, B.J. and Levitan, M.,Remote Sensing Classification of Hurricane Storm Surge Structural Damage, ASCE/SEI Structures Congress, Long Beach, 2007-04.
 72. ImageCat. ImageCat Analyzes Summer UK Floods using High Resolution Aerial and Medium Resolution Radar Imagery, http://www.imagecatinc.com/article_2007_08_a.html 2007
 73. McMillan, A., Morley, J.G., Adams, B.J. and Chesworth, S. A study of multi-temporal and multi-resolution SAR imagery for post-Katrina flood monitoring in New Orleans, International Journal of Remote Sensing. forthcoming 2007
 74. Oates, B., McMillan, A., Morley, J. and Adams, B. Empirical characterization of urban SAR response, International Journal of Remote Sensing. forthcoming 2007
 75. Womble, J. A, Adams, B. J., and Mehta, K. C. Remote-Sensing Assessment of Wind Damage," 5th International Workshop on Remote-Sensing for Disaster Response, Washington, DC,, <http://www.gwu.edu/~spi/Womble-Mehta.pdf>, 2007-09

76. Rathje, E. and Beverley J. Adams, The Role of Remote Sensing in Earthquake Science and Engineering, Opportunities and Challenges, *Earthquake Spectra*, 24(2): 471-492. 2008.