An IT Infrastructure for Responding to the Unexpected

Magda El Zarki, PhD

Ramesh Rao, PhD

Sharad Mehrotra, PhD

Nalini Venkatasubramanian, PhD

Proposal ID: 0403433

University of California, Irvine

University of California, San Diego

June 1st, 2006

Table of Contents

AN IT INFRASTRUCTURE FOR RESPONDING TO THE UNEXPECTED

Executive Summary

The University of California, Irvine (UCI) and the University of California, San Diego (UCSD) received NSF Institutional Infrastructure Award 0403433 under NSF Program 2885 CISE Research Infrastructure. This award is a five year continuing grant and the following report is the Year Two Annual Report.

The NSF funds from year two ($405, 210) were split between UCI and UCSD with half going to each institution. The funds were used to begin creation of the campus-level research information technology infrastructure known as Responsphere at the UCI campus as well as beginning the creation of mobile command infrastructure at UCSD. The results from year two include 113 research papers published in fulfillment of our academic mission. A number of drills were conducted either in the Responsphere infrastructure or equipped with Responsphere equipment in fulfillment of our community outreach mission. Additionally, we have made many contacts with the First Responder community and have opened our infrastructure to their input and advice. Finally, as part of our education mission, we have used the infrastructure equipment to teach or facilitate a number of graduate and undergraduate courses at UCI including:
UCI ICS 214A, UCI ICS 214B, UCI ICS 215, UCI ICS 203A, UCI ICS 278, UCI ICS 199, UCI ICS 290, UCI ICS 280, UCI ICS 299.

The following UCSD courses have either utilized Responsphere infrastructure, or in some cases, project-based courses have either contributed to infrastructure improvements or built new components for the infrastructure: *ECE 191 (3 projects), ECE 291 (2 projects), MAE 156 (1 project), ECE 158.*

Year two was an excellent year for Responsphere and industry relationship building. A number of strategic alliances were forged at UCI with companies such as Ether2, Printronix, CDWG, AMD, Sun Microsystems, IBM, and Psion Teklogic. Printronix generously donated most of the RFID equipment utilized in the UCI Responsphere space. AMD provided 4 dual-core (8 processor) Opteron MP 875 processors and Responsphere purchased the main board, memory, and other server components. IBM donated 22 e-series servers that are currently being used as a Beowulf cluster serving the needs of many research projects. Additionally, partnership discussions and/or non-disclosure agreement negotiations continue with Boeing, Apani Networks, Asvaco, IBM, Convera, Conexant, Vital Data Technologies, Motorola, and other information technology companies. Funding opportunities were pursued with The School Broadcasting Company in the form of a UC Discovery grant. In addition, conversations with SBC started to form a non-profit nation wide consortium on disaster warning systems.

At UCSD, we worked with several companies in preparation for the Gaslamp Quarter testbed deployment. Specifically, in the Mardi Gras Event the UCSD team participated in on February 28, 2006, Tropos Networks discounted our purchase of several access points. These access points were used as the core wireless infrastructure of the downtown area. SkyRiver Communications provided the microwave backhaul link; one of the Tropos nodes was configured as a gateway node and connected to this backhaul. Mushroom Networks, a start-up company that was founded within Calit2, provided redundancy to the Tropos and CalMesh access points and was available to replace a failure of the backhaul. UCSD has also signed a non-disclosure agreement with NetPhone, a VOIP platform venture that is also interested in the area of restoring communications infrastructure in the event of a disaster. We have been in discussions with them about mesh networking platforms, how to expand scale and range of mesh networking, as well as developing a concept command vehicle that would help with infrastructure deployment and communication restoration in a disaster.

First Responder partnerships have been essential to the success of the Responsphere project. We have formed partnerships with the Environmental Health and Safety at UCI and they have provided us with many first responder contacts as well as users of the Responsphere infrastructure. The Orange County Fire Authority provides guidance to the project with regard to first responder needs with fire department communications and data. We are working with Office of Emergency Preparedness in Los Angeles. Finally, the City of Ontario and the Ontario Fire department are working with the Responsphere project to create a disaster information portal.

Collaboration with UCSD Campus Police to validate new technologies, learn more about their needs and gain exposure to other technology-related groups within the local San Diego first-responder community have continued. Additionally, UCSD has strengthened the relationship with UCSD Emergency Management and Environment, Health and Safety (EH&S) groups. The fire chief of Foster City, CA, visited Calit2 in December 2005, and through him we were able to obtain a fire jacket to build a prototype wearable CalMesh node, and discuss the possibility of participating in one of their drills sometime in 2007.

UCSD researchers collaborated with the San Diego Police department to deploy a wireless network equipped with sensors in downtown San Diego for the February 28, 2006 Mardi Gras event. The objective was to better understand the technology requirements of the police and to see how we can help them enhance their situational awareness of crowded venues. As part of this effort, members of the Responsphere project have attended a number of emergency response events including symposiums, drills, city-wide events and workshops. At UCSD, we instrumented a number of drills with Responsphere technology and instrumented portions of the City of San Diego for the police department during the city's Mardi Gras celebration. Social science researchers from UCI participated in this event.

In addition, we have conducted a number of successful drills within the UCI infrastructure testing IT solutions and capturing data that was used to calibrate our evacuation simulator, used for First Responder training, as well as populating our data repository. For instance, the June 22, 2006 evacuation drill at the UCI campus resulted from the strategic partnership between EH&S and Responsphere. This was the first time an entire zone (Zone 3 – all 300 series buildings on map below) has been evacuated at UCI. The drill produced a number of data sets consisting of video, audio, gas levels, temperature data, people counts, and network metrics. Additionally, the drill served to validate and test a number of Responsphere technologies such as the EvacPack, the mobile cameras, the FRS Radios, and the mobile sensing platform. The social science researchers utilized the Responsphere infrastructure to conduct two experiments during the drill as well. In addition, the data sets collected during the drills are being made available to variety of researchers across various academic (e.g., Georgia Tech) including international partners such as National University of Singapore.

The CalMesh infrastructure developed at UCSD was used to provide connectivity for all of the devices used in the WIISARD project. Responsphere researchers participated in and deployed CalMesh in a number of RESCUE and WIISARD project activities. Besides UCI drills, our team also participated in campus drill at UCSD which was a simulated earthquake scenario on November 15, 2005. Also on November 15, 2005, the CalMesh team, in conjunction with the WIISARD project, participated in a drill organized by the San Diego regional Metropolitan Medical Strike Team (MMST) in Del Mar, CA. Most recently, the CalMesh team deployed their network in conjunction with WIISARD at a California Disaster Medical Assistance Team (DMAT) exercise – the "Rough and Ready" Drill in San Jose, Ca on May 5-6, 2006.

Both UCI and UCSD are currently preparing for large scale exercises in the near future. At UCSD, plans are to participate in a drill being organized by MMST, in conjunction with UCSD disaster response projects and UCSD emergency management /campus police. At UCI, we are working with EH&S and Orange County Fire Authority (OCFA) on a drill involving Bio-hazard scenario. In preparation for this drill we have acquired a Dreager Multi-gas sensor which along with variety of other sensors that are part of the Responsphere infrastructure will be deployed during the drill for environmental sensing. Furthermore, autonomous mobile sensing platforms have been developed for technology deployment during this exercise. Both UCI and UCSD drills are planned in August, 2006.

Responsphere Autonomous Mobile Sensing Platforms

Spending Plan

Spending plans for year 3 at UCI include: continuation/extension of the pervasive infrastructure (the "Smart Space"), continuing our outdoor infrastructure into Phase 3 of the build-out, increasing storage, adding computation, and enhancing the visualization cluster. Additionally, we will host a number of drills, exercises and evacuations in the Responsphere infrastructure.

Spending plans for next year at UCSD include: improvement and optimization of the Gaslamp Quarter infrastructure, optimization of CalMesh, completion of the portable visualization display, and outfitting the Wireless Communications Command and Control Vehicle trailer. Like the UCI team, UCSD will also be deploying mobile infrastructure and be participating in a variety of drills and emergency response exercises.

Infrastructure

Responsphere is the hardware and software infrastructure for the Responding to Crisis and Unexpected Events (ResCUE) NSF-funded project. The vision for Responsphere is to instrument selected buildings and an approximate one third section of the UCI campus (see map below) with a number of sensing modalities. In addition to these sensing technologies, the researchers have instrumented this space with pervasive IEEE 802.11a/b/g Wi-Fi and IEEE 802.3 to selected sensors. They have termed this instrumented space the "UCI Smart-Space."

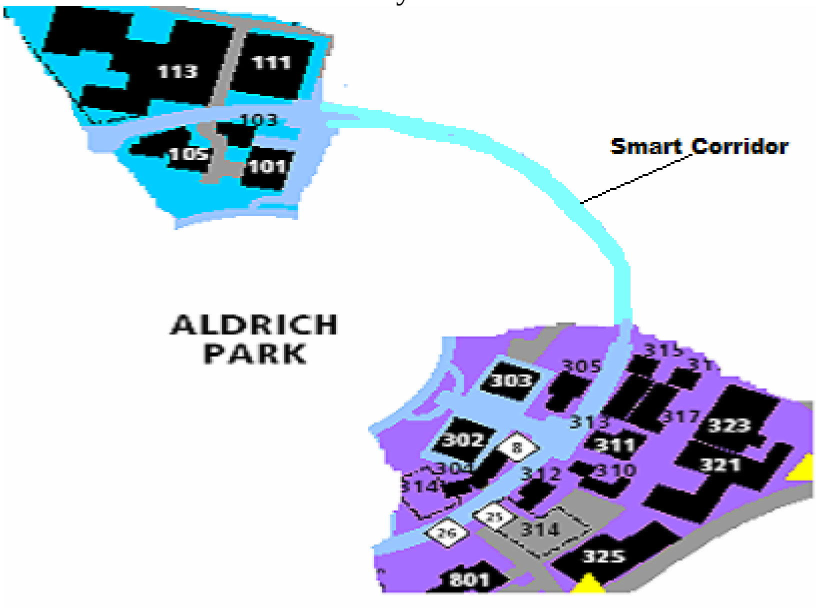

UCI Smart-Space

The sensing modalities within the Smart-Space include audio, video, powerline networking, motion detectors, RFID, and *people counting* (ingress and egress) technologies. The video technology consists of a number of fixed Linksys WVC54G cameras (streaming audio as well as video), mobile Linksys WVC 200 tilt/pan/zoom cameras, and several Canon VB-C50 tilt/pan/zoom cameras. These sensors communicate with an 8-processor (3Ghz) IBM e445 server as well as an 8-processor (4 dual-cores) AMD Opteron MP 875 server. Data from the sensors is stored on an attached IBM EXP 400 with a 4TB RAID5EE storage array. This data is utilized to provide emergency response plan calibration, perform information technology research, as well as feeding into our Evacuation and Drill Simulator (DrillSim). The data is also provided to other disaster response researchers through a Responsphere Affiliates program and web portal. Back-ups of the data are conducted over the network to Buffalo Terrastation units as well as a third generation stored off-site.

This budget cycle (2005-2006), the UCI Smart-Space Infrastructure was extended to the rest of the 300-series buildings (see map above) by purchasing 5G Wireless Extended Range Wi-Fi equipment. This equipment was installed on building 303 (see map below), and provides Wi-Fi coverage for all of the 300-series buildings. We anticipate having the entire UCI Smart-Space instrumented by budget year three.

UCSD is developing a fixed infrastructure at the downtown San Diego Gaslamp Quarter, and a mobile infrastructure which will ultimately consist of CalMesh nodes, a portable visualization display, and a mobile communications vehicle. In the Gaslamp Quarter, a core infrastructure consisting of Tropos wireless access points (with one access point configured as a gateway node and connected to a 3Mbps backhaul provided by SkyRiver Communications) was deployed in January 2006. In order to provide adequate network coverage for the Mardi Gras event (February 28, 2006), this infrastructure was supplemented by a deployment of CalMesh nodes, wireless ad-hoc mesh network gateways/routers developed at UCSD. CalMesh was used to extend the wireless coverage downtown so that it encompassed the entire area of the Mardi

Gras event. A satellite dish to provide additional backhaul capability was also deployed in the southern part of the event area.

CalMesh nodes also provided a secure VPN for a group of wireless pan-tilt-zoom cameras that were deployed at key areas of the event. Camera locations were determined with input from the San Diego Police Critical Incident Management and Special Events Units, who were responsible for the event. The cameras were installed to enable the police to gain better situational awareness of the event area and to identify specific areas (e.g., heavy crowds, public safety incidents) to which they could re-direct police officers. Due to privacy concerns, the cameras were set up the day of the event and were taken down the day after the event. Only the police officers in the Mardi Gras command center were able to view the streams as allowed by the secure VPN. UCSD demonstrated our own incident command capability by deploying an incident command center in a room across from that of the San Diego Police's incident command center.

CalMesh nodes also provide a mobile wireless ad-hoc mesh networking infrastructure to support both research and activities (training exercises and drills) for both the RESCUE and WIISARD (Wireless Internet Information Systems for Medical Response in Disasters) projects.

We have also performed mesh network compatibility testing with San Diego Engine Co. No. 35 with Captain Cherry of B Division. Electromagnetic field measurements were made in and around an operational fire truck. We measured interference between the fire truck radio system, fireman handheld radios and UCSD's mesh network wireless devices. Further measurements were conducted on interference between the fire truck's magnetic braking system and UCSD's mesh network wireless devices.

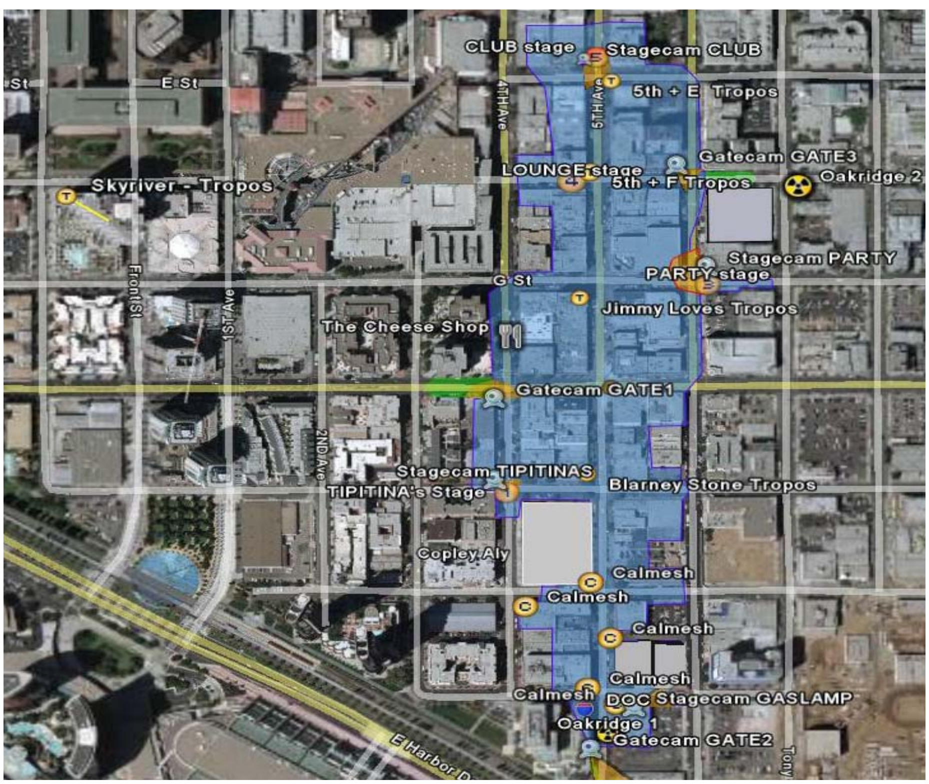

Plan of Gaslamp Quarter Infrastructure at Mardi Gras – San Diego, CA

Outreach

In fulfillment of the outreach mission of the Responsphere project, one of the goals of the researchers at the project is to open this infrastructure to the first responder community, the larger academic community including K-12, and the solutions provider community. The researchers' desire is to provide an infrastructure that can test emergency response technology and provide metrics such as evacuation time, casualty information, and behavioral models. These metrics provided by this test-bed can be utilized to provide a quantitative assessment of information technology effectiveness. Printronix, IBM, and Ether2 are examples of companies that have donated equipment in exchange for testing within the Responsphere testbed.

One of the ways that the Responsphere project has opened the infrastructure to the disaster response community is through the creation of a Web portal. On the www.responsphere.org website there is a portal for the community. This portal provides access to data sets, computational resources and storage resources for disaster response researchers, contingent upon their complying with our IRB-approved access protocols. IRB has approved our protocol under Expedited Review (minimal risk) and assigned our research the number HS# 2005-4395.

At UCI we have been active in outreach efforts with the academic community, organizing the following conferences and workshops:

1. We organized the first Earthquake Information Dissemination Workshop which brought together multidisciplinary researchers and practitioners from academia, industry, and government. The gathering included experts from the areas of seismology, computer science, earthquake engineering, social science, disaster science as well as stake-holders such as the State of California, Office of Emergency Services, local and county school officials, and other government partners at city and county levels. The team (which included members of the original Trinet study in the state of California) is moving towards enabling a real-time seismic alert and warning system for the State of California.
2. We organized a major international conference– IEEE Intelligence and Security Informatics 2006 working with our colleagues at University of Arizona, and University of Texas at Dallas. The conference attracted over 150 participants with common interest in terrorism informatics and emergency response.
3. We hosted the NSF Cybertrust PI meeting at Orange County.

Other outreach activities at UCI include: demonstrations and recruitment activities for Girls Inc., demonstrations and recruitment activities for The American Indian Summer Institute in Computer Science, demonstrations and recruitment activities for The Sally Ride Science Festival, demonstrations and recruitment activities for local area high schools coordinated through UCI Student Affairs office. Additionally, we have hosted three intern students from Telecom Lille in France as well as a number of undergraduate students including three CalIT2 Surf-IT fellows. We also are hosting one high school intern as a result of an industry affiliation with Nextgenix Inc.

At UCSD we have been active in outreach efforts with the academic community, organizing the following conferences and workshops:

1. DeFanti, Tom. Co-chair, iGrid 2005, Calit2, UCSD, September 26-30, 2005.
2. B.S. Manoj and Hubenko, Alexandra. Co-Chairs, Workshop on Future Communication Requirements for Emergency Response. Third International Conference on Information Systems for Crisis Response and Management (ISCRAM '06).

3. B.S. Manoj. General Co-Chair, The First IEEE International Workshop on Next Generation Wireless Networks 2005 (IEEE WoNGeN '05). Held in conjunction with IEEE International Conference on High Performance Computing 2005 (IEEE HiPC '05) Goa, India, December 18-21, 2005.

Responsphere researchers and technologists from both campuses 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.. We list a sample of such talks below.

1. DeFanti, Tom. Keynote speaker, IEEE International Symposium on Multimedia in Irvine, CA, December 12-14, 2005.
2. Schulze, Jurgen. "Creating Virtual Environments." Guest lecture in CSE168: Rendering Algorithms. Department of Computer Science and Engineering, University of California San Diego, June 5, 2006
3. Alexandra Hubenko: Invited talk, “CalMesh: A Wireless Ad Hoc Mesh Network for Disaster Response”. Microsoft/IAFC Fire Service Technology Symposium, Redmond, WA, December 6-7, 2005. This presentation has launched collaborations between UCI and Orange County, CA Fire Dept, and UCSD and Foster City, CA Fire Dept.
4. Researcher John Miller demonstrated ZigZag to a group of interested blind individuals on at the National Federation for the Blind (NFB) conference July 2, 2005 in Louisville, Kentucky. Zigzag is a sense of touch guiding system including a transmitter and a handheld guiding device. Miller will be presenting his updated research (ZigZag 2) at the NFB conference in July 2006.
5. Raheleh Dilmaghani: “Performance Evaluation of RescueMesh: A Metro-Scale Hybrid Wireless Network” (R. B. Dilmaghani , B. S. Manoj, B. Jafarian, R. R. Rao. WiMesh-2005: First IEEE Workshop on Wireless Mesh Networks. Held in conjunction with SECON-2005 Santa Clara, CA, September 26, 2005 (This presentation has led to a partnership with a start-up company concentrating on ad-hoc mesh network deployments for emergency response).
6. Ramesh Rao, The First IEEE International Workshop on Next Generation Wireless Networks 2005 (IEEE WoNGeN '05) - held in conjunction with IEEE International Conference on High Performance Computing 2005 (IEEE HiPC '05) Goa, India, December 18-21, 2005. BS Manoj: workshop Co-Chair; Dr Ramesh Rao: Keynote talk, “Responding to Crises and the Unexpected”
7. Information Theory & Applications Center Inaugural Workshop, Calit2, UC San Diego, February 6-10, 2006
   - BS Manoj, Chair, Session on Sensor Networks
   - Bhaskar Rao, Chair, Session on MIMO
   - Serge Belongie, Chair, Session on Bioinformatics International
8. Prof. Ramesh. R. Rao delivered the keynote talk titled “Responding to the Crises and Unexpected” during the opening ceremony of the International workshop on Next Generation Wireless Networks 2005 (WoNGeN’05).
9. Raheleh B. Dilmaghani presented her paper on “Evaluation of a Metro-Scale Wireless Mesh Network“ in IEEE Workshop on Wireless Mesh Networks (WiMesh’05), held in conjunction with SECON-2005, Santa Clara, California, 26th September, 2005.

10. Raheleh B. Dilmaghani presented her paper on “Emergency Communication Challenges and Privacy“ at the International Conference on Information Systems for Crises Response and Management (ISCRAM 2006) [www.iscram.org] held during May 14th to May 18th at New Jersey Institute of Technology, Newark, NJ.
11. Rajesh Mishra presented a paper on “Challenges in Using Distributed Wireless Mesh Networks in Emergency Response“ at the International Conference on Information Systems for Crises Response and Management (ISCRAM 2006) [www.iscram.org] held during May 14th to May 18th at New Jersey Institute of Technology, Newark, NJ.
12. Dr. B. S. Manoj presented a seminar titled “On the Evolution of MAC Protocols for Ad hoc Wireless Networks,” on 10th December 2005, at Amrita University, Amritapuri Campus, Kerala, India. Amrita University is a partner in UCSD CalIT2’s international collaboration on education and emergency response.
13. Raheleh B. Dilmaghani presented her paper on “Designing Communication Networks for Emergency Situations“ at the International Symposium on Technology and Society (ISTAS '06) during June 9-10 at New York City, NY.
14. Dr. B. S. Manoj delivered another seminar titled “On Using Multihop Wireless Relaying in Next Generation Wireless Networks,” on 12th December 2005 at the Amrita University, Amritapuri Campus, Kerala, India.
15. Dr Ramesh Rao participated in “Strengthening the Scientific and Technical Responses to Hurricane Katrina: A Meeting of Experts”, convened at the National Academy of Sciences, Washington, DC, November 14-15, 2005.
16. Dr. Ramesh Rao chairs a study at the National Academies Computer Science and Telecommunications Board titled “Committee on Using Information Technology to Enhance Disaster Management.” This committee is studying the requirements for enhancement of crisis response, and will ultimately produce a report on how information technology can enhance crisis preparedness, response, and consequence management of natural and man-made disasters. The 18-month long study continues through October 2006.
17. Naveen Ashish, SAMI project leader is a featured speaker at the Institute for Defense and Government Analysis (IDGA) seminar on Joint Search and Rescue, July 25-26, Arlington, VA. He will provide a tutorial on situational awareness technologies being investigated and developed by SAMI.
18. Naveen Ashish is co-chair for the AAAI workshop on Event Extraction and Synthesis being held the National Conference on Artificial Intelligence (AAAI) 2006 in Boston, July 17th 2006.
19. Butts, Carter T. Invited talk: ``Likelihood-based Network Comparison Using Permutation Models." Invited Workshop Tutorial, NetSci International Conference and Workshop on Network Science. Bloomington, IN. (5/2006).
20. Petrescu-Prahova, Miruna and Butts, Carter T. Conference presentation: ``Network Inference with Missing Data: A Performance Comparison of Existing Models." (4/2006). 26th Sunbelt Network Conference (INSNA), Vancouver, BC.
21. Tsudik, Gene. Invited talk: “Next Generation Internet Security: The Clean-Slate Approach,” NSF Invitational workshop, CMU, August 2005
22. Tsudik, Gene. Panelist “Whither RFID Security?” IEEE Securecomm 2005, Athens, September 2005

23. Tsudik, Gene. Invited talk: "The Future of Internet Security," National Research Council (NRC/NAS) study group on Power Grids, April 2006.

UCSD K-12 outreach activities include sponsoring a total of 4 student 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.

Responsphere Drills

- *November 15, 2005:* UCSD Campus Exercise – RESCUE teams participated in this campus exercise, with researchers deployed both inside the emergency operations center (EOC) and in the field. Working with campus police, emergency management, and the HazMat team, researchers deployed several technologies, including video cameras to film building evacuations, a microphone array to record inside the EOC, and participated in general field observations to better understand response situations.
- *November 15, 2005:* MMST Drill at Del Mar Fairgrounds: http://www.calit2.net/newsroom/article.php?id=745
- *February 28, 2006:* Mardi Gras in the Gaslamp Quarter, San Diego: http://www.calit2.net/newsroom/article.php?id=810
- *June 22, 2006*: UCI Campus Zone 3 evacuation drill (1/5 of the UCI campus): - RESCUE teams participated in this exercise in collaboration with the UCI campus Environmental Health and Safety office. This drill helped RESCUE researchers to study crowd dynamics, and the opportunity to test camera instrumentation, mobile 802.11 cameras, Evacpack, and DrillSIM.
- *August 22, 2006 (planned)* MMST Drill at Calit2/UCSD – RESCUE, WIISARD, Responsphere project teams.
- *August 29, 2006 (planned)* BioHazard Drill at UCI, RESCUE and Responsphere teams.

Research Demonstrations with First-Responder, Government, and State Community Groups:

- October 10, 2005. A RESCUE demonstration was presented to the City of Ontario Emergency Services group, resulting in a partnership and development of the Ontario Portal (see artifacts section for details)
- August 17, 2006 ZigZag sense of touch guidance system trials (including participation /input of UCSD HazMat team)
- January 10, 2006: RESCUE All Hands Meeting Community poster session and demonstration exposition: Participants included members of CAB, TAC, and community partners (UCSD Police, Emergency Management/EHS, and other disaster response community partners)
- In January 2006, RESCUE researchers presented the PaDOC Demo (Privacy-Preserving Video Surveillance) to Peter Freeman of NSF

- February 9, 2006: RESCUE researchers presented a demo to representatives from the California Governor's Office – David Crane, Special Advisor of Jobs and Economic Growth, and Mindy Fletcher, Deputy Chief of Staff for External Affairs.
- Calit2/UCSD-TUM Automotive Software Workshop, March 15-17, 2006, UCSD
- Demonstrations for San Francisco Mayor Gavin Newsom and staff, March 23, 2006

Responsphere Management

The Responsphere project leverages the existing management staff of the affiliated RESCUE project which is a NSF funded Large ITR. In addition, Responsphere, given the scale of the technology acquisition and deployment has hired technologists who are responsible for purchase, deployment, and management of the infrastructure. The management staff consists of a Technology Manager (Chris Davison), Project Coordinator (Jean Chin – leveraged from RESCUE), and Project Manager (Quent Cassen – leveraged from RESCUE). At UCSD, the management staff consists of a Project Manager (Alex Hubenko) and Project Support Coordinator (Helena Bristow). The management staff and technologists associated with Responsphere possess the necessary technical and managerial skills for both creation of the infrastructure and collaboration with the industry partners. The skill set of the team includes: Network Management, Technology Management, VLSI design, and cellular communications. This skill set is crucial to the design, specification, purchasing, deployment, and management of the Responsphere infrastructure.

Part of the executive-level decision making involved with accessing the open infrastructure of Responsphere (discussed in the Infrastructure portion of this report) is the specification of access protocols. Responsphere management has decided on a 3-tiered approach to accessing the services provided to the first responder community as well as the disaster response and recovery researchers.

Tier 1 access to Responsphere involves a read-only access to the data sets as well as limited access to the drills, software and hardware components. To request Tier 1 access, the protocol is to submit the request, via www.responsphere.org, and await approval from the Responsphere staff as well as the IRB in the case of federally funded research. Typically, this access is for industry affiliates and government partners under the supervision of Responsphere management.

Tier 2 access to Responsphere is reserved for staff and researchers specifically assigned to the ResCUE and Responsphere grant. This access, covered by the affiliated Institution's IRB, is more general in that hardware, software, as well as storage capacity can be utilized for research. This level of access typically will have read/write access to the data sets, participation or instantiation of drills, and configuration rights to most equipment. The protocol to obtain Tier 2 access begins with a written request on behalf of the requestor. Next, approval must be granted by the Responsphere team and, if applicable, by the responsible IRB.

Tier 3 access to Responsphere is reserved for Responsphere technical management and support. This is typically "root" or "administrator" access on the hardware. Drill designers could have Tier 3 access in some cases. The Tier 3 access protocol requires that all Tier 3 personnel be UCI or UCSD employees and cleared through the local IRB.

Personnel

University of California Irvine (UCI)

Name	***Role(s)***	***Institution***
Naveen Ashish	Visiting Assistant Project Scientist	UCI
Carter Butts	Assistant Professor of Sociology and the Institute for Mathematical Behavioral Sciences	UCI
Howard Chung	ImageCat	Inc.
Remy Cross	Graduate Student	UCI
Mahesh Datt	Graduate Student	UCI
Rina Dechter	Professor	UCI
Mayur Deshpande	Graduate Student	UCI
Ronald Eguchi	President and CEO	ImageCat
Magda El Zarki	Professor of Computer Science	UCI
Ramaswamy Hariharan	Graduate Student	UCI
Bijit Hore	Graduate Student	UCI
John Hutchins	Graduate Student	UCI
Charles Huyck	Senior Vice President	ImageCat
Ramesh Jain	Bren Professor of Information and Computer Science	UCI
Aaron Jow	Graduate Student	UCSD
Dmitri Kalashnikov	Post-Doctoral Researcher	UCI
Iosif Lazaridis	Graduate Student	UCI
Chen Li	Assistant Professor of Information and Computer Science	UCI
Yiming Ma	Graduate Student	UCI
Gloria Mark	Associate Professor of Information and Computer Science	UCI
Daniel Massaguer	Graduate Student	UCI
Sharad Mehrotra	RESCUE Project Director, Professor of Information and Computer Science	UCI
Shivajit Mohapatra	Graduate Student	UCI
Miruna Petrescu-Prahova	Graduate Student	UCI
Vinayak Ram	Graduate Student	UCI
Will Recker	Professor of Civil and Environmental Engineering, Advanced Power and Energy Program	UCI
Nitesh Saxena	Graduate Student	UCI
Dawit Seid	Graduate Student	UCI
Masanobu Shinozuka	Chair and Distinguished Professor of Civil and Environmental Engineering	UCI
Michal Shmueli-Scheuer	Graduate Student	UCI
Padhraic Smyth	Professor of Information and Computer Science	UCI
Jeanette Sutton	Natural Hazards Research and Applications Information Center	University of Colorado at Boulder
Nalini Venkatasubramanian	Associate Professor of Information and Computer Science	UCI

Kathleen Tierney	Professor of Sociology	University of Colorado at Boulder
Jehan Wickramasuriya	Graduate Student	UCI
Xingbo Yu	Graduate Student	UCI

University of California San Diego (UCSD)

Name	***Role(s)***	***Institution***
Ramesh Rao	PI; Professor, ECE; Director, Calit2 UCSD Division	Calit2, UCSD
John Miller	Senior Development Engineer	Calit2, UCSD
Ganapathy Chockalingam	Principal Development Engineer	Calit2, UCSD
Babak Jafarian	Senior Development Engineer	Calit2, UCSD
John Zhu	Senior Development Engineer	Calit2, UCSD
BS Manoj	Post-doctoral Researcher	Calit2, UCSD
Sangho Park	Post-doctoral Researcher	Calit2, UCSD
Stephen Pasco	Senior Development Engineer	Calit2, UCSD
Helena Bristow	Project Support	Calit2, UCSD
Alexandra Hubenko	Project Manager	Calit2, UCSD
Raheleh Dilmaghani	Graduate Student	ECE, UCSD
Shankar Shivappa	Graduate Student	ECE, UCSD
Wenyi Zhang	Graduate Student	ECE, UCSD
Vincent Rabaud	Graduate Student	CSE, UCSD
Salih Ergut	Graduate Student	ECE, UCSD
Javier Rodriguez Molina	Hardware development engineer	Calit2, UCSD
Stephan Steinbach	Development Engineer	Calit2, UCSD
Rajesh Hegde	Postdoctoral Researcher	Calit2, UCSD
Rajesh Mishra	Senior Development Engineer	Calit2, UCSD
Brian Braunstein	Software Development Engineer	Calit2, UCSD
Mustafa Arisoylu	Graduate student	ECE, UCSD
Tom DeFanti	Senior Research Scientist	Calit2, UCSD
Greg Dawe,	Principal Development Engineer	Calit2, UCSD
Greg Hidley	Chief Infrastructure Officer	Calit2, UCSD
Doug Palmer	Principal Development Engineer	Calit2, UCSD
Don Kimball	Principal Development Engineer	Calit2, UCSD
Leslie Lenert	Associate Director for Medical Informatics, Calit2 UCSD Division; Professor of Medicine, UCSD; PI, WIISARD project	Calit2, UCSD
Troy Trimble	Graduate Student	ECE, UCSD
Cuong Vu	Senior Research Associate	Calit2, UCSD
Boz Kamyabi	Senior Development Engineer	Calit2, UCSD
Jurgen Schulze	Postdoctoral Researcher	Calit2, UCSD
Qian Liu	Systems Integrator	Calit2, UCSD
Joe Keefe	Network Technician	Calit2, UCSD
Brian Dunne	Network Technician	Calit2, UCSD
Cuong Vu	Senior Research Associate	Calit2, UCSD

Responsphere Research Thrusts

The Responsphere Project provides the IT infrastructure for Rescue project. The project is divided into the following five research projects: Situational Awareness from Multi-Modal

Input (SAMI), Policy-driven Information Sharing Architecture (PISA), Customized Dissemination in the Large, Privacy Implications of Technology, and Robust Networking and Information Collection. The following research and research papers (by project area) were facilitated by the Responsphere Infrastructure, or utilized the Responsphere equipment.

Situational Awareness from Multimodal Inputs (SAMI)

The SAMI project has taken on the challenge, within RESCUE, of developing technologies that dramatically improve situational awareness for first-responders, response organizations, and the general public. This translates to the following scientific goals: (1) develop an event-oriented situational data management system that seamlessly represents activities (their spatial, temporal properties, associated entities, and events) and supports languages/mechanisms/tools to build situational awareness applications, (2) create a robust approach to signal analysis, interpretation, and synthesis of situational information based on event abstraction. Our goals also include the development of two artifacts -- an information reconnaissance system for disaster data ingest, and an integrated situational information dashboard that aids decision making.

Activities and Findings

Presented below is a summary of progress to date in each of the three SAMI research areas: situational data management; signal interpretation, analysis and synthesis; and analyst tools. Special attention is given to highlighting innovative research contributions.

Situational Data Management

This component is concerned with the development of a data management system for situational data, i.e., a system that provides storage, modeling and representation, and querying capabilities over situational data in much the same way that a relational database system provides such capabilities for structured enterprise data. The progress in the research thrusts for this component is described below.

Situational data modeling and representation
The team working on event modeling has been focusing on the development of data model and management tools for situational awareness. Two challenges – in particular - have been addressed:

Challenge 1 (Multimodal Event Model and Data Store): For situational awareness, the ability to access information about ongoing events and to analyze these events is crucial. We are working on a common event store that permits a unified storage of events detected from arbitrary sources and any kind of documenting media available for them. There are several challenges involved in this work. The first is the development of a uniform, generic event model that permits the representation and interrelation of events of different types, detected from different sources and modalities, and addressing different levels of abstraction down from sensor data-level events up to domain-level events. Such a model must address the uncertainty of each event.

Secondly, we must provide an event ontology language that allows a basic event to be described. For instance, time, location, structure, causality, and documenting media, need to be represented for a) different types of events from different sources and b) different levels of abstraction. This will ultimately permit the specification of event inference rules for the detection of events. We must also develop an event retrieval language that considers uncertainty, as well as implement an event store based on the developed event model, ontology language, and retrieval language.

With regard to Challenge 1, we report the following research findings and contributions. First, for the E multimedia event model, E is generic and offers expressive primitives for capturing the basic aspects of event description for any kind of event from any source and at any level of abstraction. These include freely definable event types, properties, tags, and relationships between events. E permits one to capture uncertainty at the level of events as well as at the level of individual event properties. E allows events to refer to all documenting media and sensor data available on them, thereby providing a unified index onto media and sensor data. As a first application, E has been successfully applied to the representation and analysis of outdoor reconnaissance patrol events gathered from wearable sensors and cameras. We have also implemented the EDB event store for storing events based on E. EDB provides basic event storage and navigational event access services on top of a relational database management system via a web service API. We are still in the process of defining an event ontology language and developing event retrieval algebra.

Challenge 2 (Distributed Sensor Data Acquisition and Transformation Infrastructure): In order to detect events with which the above mentioned event store is about to be filled, developers of media and sensor data analysis algorithms for event detection are in need of an infrastructure that supports them in accessing and processing the data streams produced in a multimodal sensor data network. Without such an infrastructure, developers must access each sensor data stream explicitly and deal with non-trivial issues such as different sensor communication protocols, different data stream formats, data stream synchronization, network load and failure, and optimization of parallel data stream processing tasks. Current sensor data management infrastructures and data stream management systems, however, do not provide the developers of multimodal sensor data analysis algorithms with such an infrastructure. Common problems include: a) lack of support for raw data streams such as video, b) a centralized stream processing with all incurring problems of scalability and network bottlenecks, c) lack of declarative query languages allowing one to deal with all of raw streams, and d) structured streams, or streams of events, and limited optimization capabilities.

With regard to this challenge, we report on the following research findings and contributions. First, we have implemented CAMAS VM, a distributed sensor data processing and transformation infrastructure. It features the concept of sensor data processing nodes. These can be installed on the machines of a network, take an arbitrary number of raw and structured sensor data streams as inputs, apply an importable sensor data transformation function code on the input streams, and deliver an output stream as a result. CAMAS VM nodes can be connected to form arbitrarily complex sensor data transformation topologies.

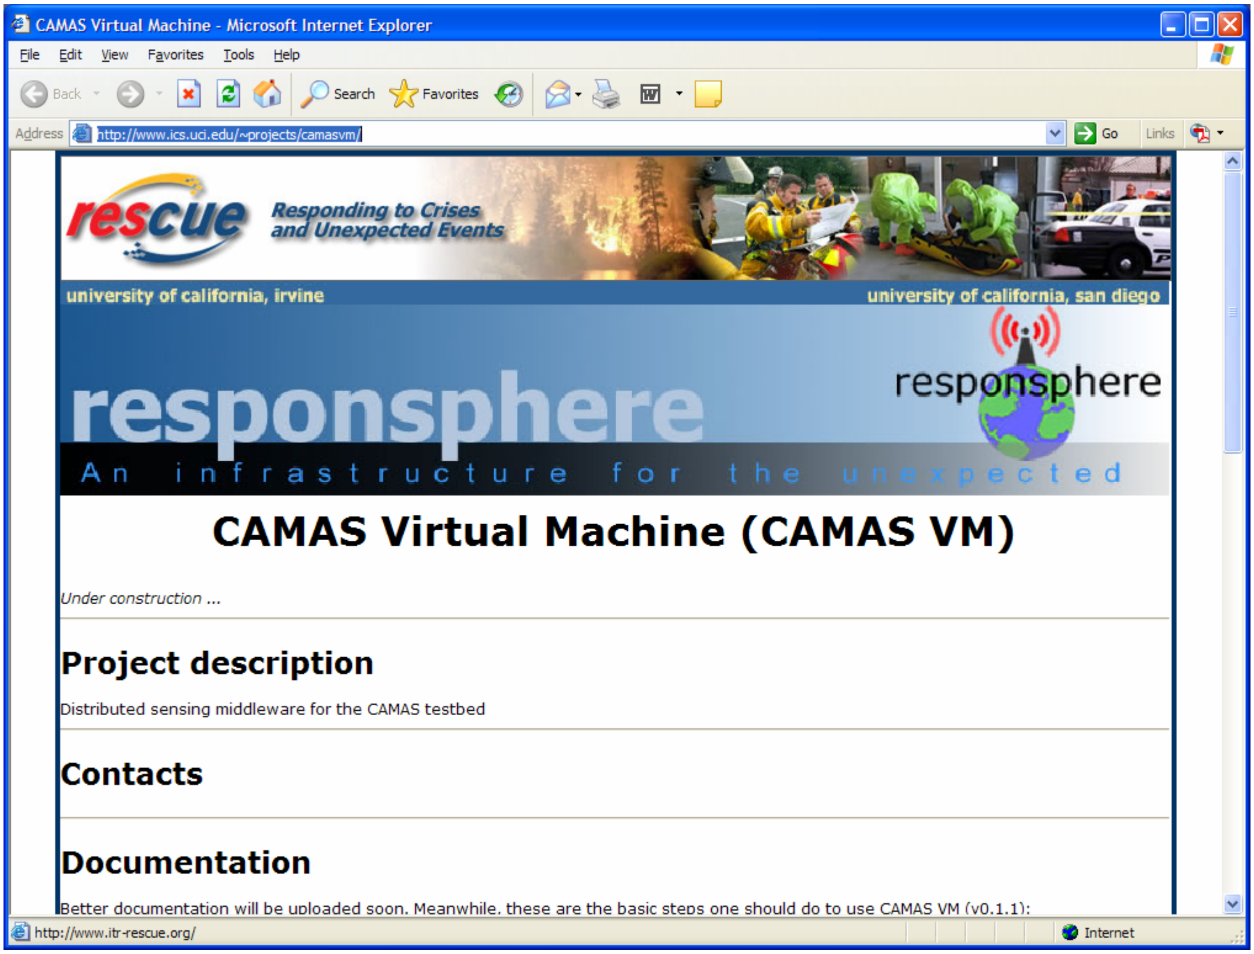

CAMAS Virtual Machine http://www.ics.uci.edu/~projects/camasvm

We have defined a basic sensor data transformation function code, covering acquisition of sensor data from video cameras and people counters and the synchronization of different sensor data streams. A registry has been implemented for CAMAS VM. The registry maintains information of the CAMAS VM nodes installed, the various types of sensors available, information about the network topology, as well as the sensor data transformation code available for injection.

Finally, we have defined and implemented a sensor data acquisition and transformation language (SATLite). SATLite allows developers to declaratively specify sensor data transformation topologies at an abstract level. These are processed by a query processor which in cooperation with the registry initiates these topologies on the nodes of the network. We have implemented a translator for converting an essential subset of CQL to SATLite. We have proven the expressiveness of our approach.

Spatial reasoning and querying

Work in the area of spatial reasoning and querying has been ongoing in RESCUE and SAMI for over a year now. We are working on multiple aspects of the problem including the extraction of spatial expressions and references from free text, the interpretation of spatial expressions in a probabilistic manner, and indexing mechanisms for efficient querying of the spatial references given complicated probabilistic representations.

Research findings and contributions: We have developed a new end-to-end situational awareness architecture that is specifically tailored towards emergency management requirements. The framework uses human reporters (first responders, concerned citizens etc.) as sensors to gain raw information and basic information on a crisis. The framework currently deals with free text inputs -- one of the major input modalities in almost all the emergency management applications.

In the past year, we have focused on the spatial reasoning and querying components that are commonly used in the applications. Reasoning or querying directly on the raw input often fails to achieve the desired retrieval quality. To attain a higher level of spatial knowledge, we have implemented a spatial information extractor based on various existing spatial Gazetteers, and represent the extracted knowledge as spatial expressions for our modeling tools. To capture the uncertainties in the spatial expressions, we have proposed a novel modeling process to convert the spatial expressions to probabilistic representations. By analyzing the emergency management query requirements, we defined formal semantics of the commonly-used spatial queries given probabilistic data representations. To speed up the query response time, we have proposed a set of novel indexing structures based on a Grid indexing framework. We have extensively tested our approach on a real emergency response dataset, and demonstrated its effectiveness and efficiency as compared to the current state-of-art.

Challenges: The primary challenges in this area at this stage are related to the imprecision in information, imprecision in the (raw) input information per se, and also imprecision and uncertainty in the information extracted from such input data in an automated fashion. Specifically, due to the limitations of automated natural language understanding, our spatial extractor can make mistakes in recognizing the spatial references and relations in a report. These mistakes will further propagate to the probabilistic event modeling stage. Given these uncertainties and errors made at extraction and modeling stages, the challenge for us is to device a formal retrieval mechanism that can cope with these uncertainty factors and errors, and still be able to retrieve the relevant reports/events in an effective way.

Situational data querying

At this point, we have completed the implementation of an event data store over a relational data management system. This data store provides basic (relational) querying capabilities. The issue of querying situational or event data, treating events as first class objects is at present best captured by our work on exploratory event analysis, described in the Analyst Tools section.

Knowledge management

Our work in the knowledge management aspect of SAMI is - so far - work that has been driven by requirements in other areas such as information extraction and information refinement. In information extraction, we are developing an approach to event extraction that exploits semantics and domain knowledge in the form of ontologies. In information refinement, we are exploiting domain ontologies in relation to analysis for web appearance disambiguation. We have defined or incorporated existing domain ontologies and have used them in extraction or disambiguation tasks.

It is likely that in the coming year, the use of semantics will be more extensive and pervasive throughout the SAMI system; for instance, exploratory event analysis is another area where the use of ontologies has been initiated. We will over the next year develop an integrated knowledge management component within SAMI that will serve the needs of these multiple uses.

In summary, we have made considerable progress in a new area which is event modeling and representation and have also made progress in the ongoing area of spatial awareness. As the work on event modeling progresses, we need to make progress on the development of situational and event data querying capabilities. Some work on knowledge management has been initiated as a result of requirements in our ongoing work on information extraction from text. We are integrating the use of domain ontologies in our work on event extraction as well as in information refinement in our work on web-page disambiguation. Throughout this year, we must formalize our use of knowledge management technology (i.e., ensure that the various sub-areas such as extraction, refinement, or even situational data querying that are or will use knowledge (in a form such as ontologies) do so in a standard, unified, and interoperable manner).

Signal Interpretation, Analysis, and Synthesis

This component in SAMI deals with the conversion of raw information in sensor, text, video and audio formats to structured, cleaned, and synthesized formats that are amenable to querying and analysis. We have work-in-progress in a variety of areas ranging from audio-visual event extraction to extraction and synthesis from text to information refinement. This work is described below:

Audio-visual event extraction and synthesis
The work in this area spans a number of thrusts in speech processing, namely multi-microphone speech processing, stream segregation, and emotion detection from speech signals; in visual information analysis namely video based event detection and mobile vision; and finally multi-modal event extraction.

Multi Microphone Speech Processing: Research in this area is focused on developing robust speech processing and recognition techniques, particularly in noisy conditions. Multi-microphone techniques provide an efficient framework for canceling out ambient stationary (car, aircraft noise) and non-stationary interference (door slams, sirens) coming from different directions and mixing eventually with the speech signal of interest. Beamforming algorithms detect the direction of arrival of the speech signal of interest and cancel out the undesired interference coming from other directions which, in turn, improves speech recognition performance significantly in noisy conditions.

The major findings and contributions include development of a robust broadband adaptive beamforming algorithm. Inherently, beamforming algorithms assume the look direction of the desired speech signal to be perfectly known but in practice there is some uncertainty. We have developed a robust broadband adaptive beamforming algorithm, which combines the robustness of the delay and sum (DS) beamforming in the look direction with the uncertainty associated with the high interference rejection capability of the conventional adaptive beamforming algorithm. In addition, we have developed a corresponding robust narrowband adaptive beamforming algorithm to address computational complexity issues. The proposed robust broadband adaptive beamforming algorithm is robust to the spatially spread source.

Other activities have included design of a handheld two-microphone array system; real data from this system has been recorded in the command center at an earthquake drill and also in the Responsphere command center in the GLQ Mardi Gras drill. Finally, we have studied the approach in the context of the real world recorded Multi-channel Overlapping Numbers Corpus (MONC).

Single Channel Auditory Stream Segregation, Speech Enhancement, and Robust Speech Recognition: Major findings and contributions to date include a scheme for detecting undesired stationary and non-stationary events. Multiple speakers have been formulated using single channel speech data, as single channel speech stream segregation and speaker separation is important in significantly improving speech recognition performance in noisy conditions.

Sinusoidal modeling is conventionally used in speech modeling and coding of a single speech source. This technique models harmonics of the speech source as three components: amplitudes, phases and frequencies. This technique has been used to model and separate speech and non speech (noise) sources as well as multiple speech sources (speakers). Sinusoidal plus residual modeling and auditory grouping has been used to separate multiple speech sources with well separated pitch. Constrained iterative sinusoidal analysis and synthesis of noisy speech has been formulated using residual interpolation and robust pitch tracking for single channel speech enhancement.

Emotion Detection from Speech Signals for Enhanced QOS in Emergency Networks: Emotion Detection from speech signals is a challenging area where the emotional state of a person (happy, sad, angry, etc.) is detected automatically from the speech signal. Our work on emotion detection is focused on automatically detecting the state of the speaker/caller (Distressed or Normal) from the speech signal, without manual intervention. We have developed emotion detection techniques using novel features extracted from the speech signal, and techniques for feature selection for pruning less discriminative features

Video-based Events Detection for Enhanced Situational Awareness: Our work is centered on “homography", which in computer vision denotes a linear projective transformation that relates two different views of a scene. Representation of objects in the “homography domain” refers to the warped representation of a given object.

Usually, three-dimensional objects violate the constraints of planar points and are warped by the homography transformation. This warped representation of objects can provide very useful information when multiple homography representations are combined together. Combining multiple homography relations from multiple cameras significantly enhances robustness in image segmentation.

Our work is based on a novel approach to track people and vehicles in crowded scenes using multiple cameras. Camera calibration is not needed in our approach to extract the world coordinates of objects. Our study shows that object representation and tracking in the homography domain, independent of detailed appearances of the objects, is an efficient and robust way of handling visual events. Semantic understanding of visual events is achieved by spatio-temporal analysis of the objects in the multiple homography representations.

We also plan to develop methodologies to combine the homography-domain and the image-domain representations to analyze and extract more useful features in understanding visual events, as the homography-based representation is a projection of objects from 3D to 2D space, and lacks the detailed image appearance of the objects. Finally, foreground-background segmentation in video is a challenging task in computer vision. We plan to investigate efficient methods to handle shadows in multi-view homography domains and to segment complete shapes of objects for robust high-level vision processing.

Mobile Vision: This activity includes the study of the development of workflow/tools to produce the mobile vision platform as well as binaries and initial research into optimized libraries/hardware to enhance platform performance.

Event extraction and synthesis from text

Our work in event extraction from text is driven by the end applications of triaging and analysis of the extracted event information. We are interested in the extraction of ‘events’ from situational text data from both corpuses of documents (such as a collection of new stories) as well as real-time event extraction from transcribed audio streams.

We have developed an architecture for an end-to-end event extraction system driven by predicate logic based extraction rules. The architecture includes the definition of individual components and their interaction, and the detailed specification of the rule language for performing the extraction. One of the unique advantages of this architecture is that the extraction rules can capture and exploit text and sentential features of various types including data or knowledge based features, features obtained from linguistic analysis, and statistical features.

Our work in event extraction and synthesis from text has led us to conduct extensive literature and system surveys on event extraction and information extraction in general. This includes obtaining familiarity in some areas quite new to us such as conversation and dialogue modeling and understanding. We have also initiated the use of the open source information extraction toolkit GATE on a number of SAMI extraction subtasks as an application development engine, as well as a platform for building advanced extraction systems. In preparation for future work, we have outlined a new area of “embellishment” in information extraction and have initiated work on Web-page disambiguation as an initiation of work in this area.

Exploitation of multi-modality and semantics for extraction and synthesis

In this area, we have successfully designed and implemented a simple system that detects fundamental events from audio-video data in a meeting room. Our emphasis has been on the real-time operation of the system (synchronizing capture and processing of audio and video data). Participation in drills and real-life events has played a major part in broadening our research. Our participation in the Mardi Gras event in Downtown San Diego (March 2006) is one example where RESCUE was able to collect audio and video information from the RESCUE command center for extraction and synthesis. These drills and events also enabled us to study system deployment issues to capture synchronized audio and video data.

Information Refinement/Disambiguation
The problem of disambiguation arises in a variety of contexts, especially when information is merged from heterogeneous sources to create a unified database. It naturally arises when reports, for example, from citizens and transcripts of communications among the first responders, about events that constitute a crisis are compiled and need to be analyzed to create an overall situational awareness. The problem that we study is the reference uncertainty problem; objects are commonly referred to in real-word data via descriptions, not unique identifiers, causing ambiguity (e.g. "Washington", "J. Smith", etc).

There are two types of disambiguation problems that are most frequently studied in the literature: fuzzy lookup and fuzzy grouping. We have developed a domain-independent solution for both of these problems which are composed of new algorithms and methodologies that are based on a new paradigm that we have proposed. They improve disambiguation quality by taking into account the relational nature of many datasets and analyzing connections of entities via relationships. For fuzzy lookup, we have developed a comprehensive and principled solution, which converts the problem into solving an optimization OR problem (nonlinear programming). We recently have proposed initial algorithms for making the lookup approach self-tunable to the dataset being processed. During the last year, we also have made several important adjustments to the base algorithms employed by the lookup solution; in particular, we have reworked some of the objective functions used by the NLP optimization problem. For fuzzy grouping, we have developed two new clustering algorithms, which take into account the information about entity and relationships. One is based on agglomerative clustering and the other one is based on correlation clustering. A method for calibrating the approach to datasets being analyzed has been proposed.

We have started to look at the problem of Web Disambiguation. Our framework is capable of analyzing word distributions over webpages, weblinks and emails. It also extracts locations, organizations, and people's name and constructs a graph out of this information, which the framework is capable of analyzing. We have also devised an initial solution for incorporating ontologies into our framework, as well as starting initial research on proper modeling of web-pages, which allows us to associate the right references with the right entities and relationships. The initial findings are quite encouraging. Several of our algorithms have been redesigned to work with the problem of Web Disambiguation; in particular we have developed a new algorithm that is based on correlation clustering.

Challenges: One of the primary challenges in disambiguation is finding good datasets for experimental analysis. First, the chosen datasets must be a real-world dataset, as they are preferred by the community. Second, such a dataset must contain a real disambiguation challenge, as the alternative method -- of introducing errors manually -- is not always preferred. Third, to compute the quality of any approach, the "ground truth" should be known for that dataset. That is, this dataset should be labeled (manually) to know the matches for references with 100% certainty. Thus, in some of the research communities the solution is to pick a small dataset and hand label it. That however is often suboptimal, as the database community strongly prefers very large datasets.

Information Reliability

This is an area where we have recently initiated some development activities. One area is to develop and demonstrate the incorporation of reliability measures in an information pipeline that seeks to address an end-user analysis query for a specific disaster analysis task. Specifically we are focusing on a use where we have a grounded end-user query related to damage assessment (such as say damage to oil rigs in a particular hurricane). One of the artifacts, the Hurricane Portal described below, will be used to gather information from news and other sources on the internet and then extract and integrate it to provide an integrated assessment of the mentions of oil rig damage. We will develop reliability measures for the information coming through this pipeline, where the information is obtained from a variety of internet information sources with varying degrees of reliability and the reports themselves having varying degrees of certainty. We have initiated the formulation of a plan in collaboration with ImageCat towards realizing this objective.

There are also other possibilities in consideration for incorporation of information reliability in SAMI. One is the development and demonstration of reliability and confidence measures for assertions made by people in an environment such as a public information portal which also allows for posting of messages. We may evaluate this in the context of the Ontario information portal. Other possibilities, in the context of applicability to information dissemination are also being discussed.

While this is an area where we have only recently focused our energies on, it remains a critical aspect of the SAMI system. It is thus imperative that we make significant progress in this area next year.

People Awareness and Predictive Modeling

There are two ongoing activities within this area. The first is the development and implementation of an approach to detecting (anomalous) events from sensor data. The aim is to identify abnormal or anomalous behavior patterns from time series sensor data. The approach has been implemented and evaluated in the specific domain of detecting anomalous events from traffic highway loop-sensor data. We have developed a new probabilistic framework that can learn to recognize abnormal behavioral patterns from time-series of sensor data related to human activity (e.g., people-counter sensors in buildings and loop-sensor data for freeway traffic). Our approach is based on a novel hidden Markov Poisson model that can be learned automatically from data, separating the predictable hourly/daily/weekly rhythms of human behavior from abnormal events. In addition, we have tested and validated the methodology on several months worth of (a) Responsphere people-counter data from the UCI Calit2 building, and (b) loop-sensor traffic data from Los Angeles freeways. The experimental results showed that our new approach can accurately detect events that are atypical (e.g., meetings in the Calit2 building or large sporting events in the evenings or on weekends in traffic data). In systematic comparisons with a more standard baseline approach (using a threshold-based Poisson statistic), our new method had average accuracies of between 90 and 100% in detecting known events compared to between 70 to 80% on average for the more traditional baseline method.

The second project in this area focuses on "people forecasting" where we are developing an approach for posterior estimation based on probabilistic network analysis. Included in this work are novel parameterized importance sampling algorithms for posterior estimation in probabilistic networks that have a large number of zero probabilities. Our sampling algorithms operate on a framework of mixed networks introduced in our previous work (see RESCUE publications list). We have also developed a novel approximate reasoning class of algorithms that integrate search with importance sampling. These algorithms are able to find better approximations of posterior beliefs than the state-of-the-art algorithms on many large benchmark instances.

Challenges: The primary challenges in the above areas at this stage relate to scalability issues (to large distributions) and also in obtaining or working with erroneous loop-sensor data which is apparently quite common. Specifically, 1) solving large scale distributions exactly to determine the quality of approximation of our sampling algorithms, 2) getting plausible loop-sensor data; one in which there is little/no erroneous data, and 3) developing a model of determining when loop sensors are generating correct data and when they are not (i.e. they are damaged).

We have also had opportunities for exploring synergies between different research thrusts. For instance, in our work on embellished information extraction, we are attempting to combine strengths from both the information extraction and disambiguation areas. Also, a considerable amount of time has been spent in formulating and crystallizing the high-level directions in this component. We are now stepping up this research activity in 2006.

Analyst Tools

The Analyst Tools component of SAMI represents the information mining, analysis and visualization capabilities that we are developing and providing over the assimilated situational data. We have progressed on a couple of fronts in this area; work on exploratory analysis of event data has been ongoing for two years within SAMI and RESCUE and the focus has been on the development of a graph-analysis language suited to events. We have also been working on issues related to search and integration of GIS data and datasets.

Graph Analysis

We have completed the design of GAL (Graph Analysis aLgebra), a semantic graph query algebra that enables a user to query text-extracted data about events and their relationships. Our work on GAL is motivated by the observation that semantic graphs provide a unified representation of data, domain semantics, knowledge, and context extracted from structured, semi-structured as well as free text sources. Our GAL algebra, which is implemented as a middleware between data sources and front-end applications incorporates a number of novel operators and semantics including graph grouping, graph aggregation, semantic joins, inference and multi-structural composition. We believe (and are working on demonstrations) that these features offer new possibilities of event analysis and mining applications. The preliminary performance experiments we have conducted on GAL have provided positive results generally showing that the rate of growth of query runtimes is almost linear with the size of the database.

Included in this design was the development of various optimization techniques for efficient execution of GAL. Since we implemented GAL by translating it to SQL, the optimization techniques are aimed at two aspects of the translation: 1) achieving the generation of efficient SQL, and 2) for the operations (like sub-graph extraction) which the relational engine evaluates using inefficient plans, we have developed better techniques for plan costing and join ordering.

Challenges: The challenges in this area currently center on assimilating datasets where events and event relationships have been extracted from the raw data. Extraction (from text) is, in general, a difficult task, however, we hope that some of the issues associated with this topic can be addressed by our research on information extraction from text.

GIS querying and visualization

The second thrust focuses on GIS data, specifically searching (hidden) GIS data on the Web. We are developing ways to more effectively collect, represent and index the, often “hidden” or deep-Web GIS datasets present on the Web, so that search for such datasets becomes more effective. Current technologies make use of high-level metadata, compiled manually, for searching certain published GIS sources. However, this mechanism has the limitations of limited understanding of how the real GIS data looks and what exactly or approximately is contained in the dataset.

Our solution to these limitations is to develop a novel GIS data integration technique that incorporates the spatial analysis of query rectangle, thereby combining the best among top ranked sources. The main reasoning for this is that a list of top-k data sources for a given spatial query is not enough as each of the data sources may partially satisfy the query. Therefore, the algorithm incorporates space partitioning techniques to incrementally improve the quality of answers to spatial queries.

As part of the research process for searching GIS data, we have identified and collected publicly available GIS data sources which include around 7500 GIS layers from about 350 servers. These datasets have been cleaned, getting rid of the original schema, and incorporated into a uniform schema which contains only geography and keywords from each dataset. We are now in the process of developing a framework to represent and index terabytes of GIS data so that an efficient search can be done on such datasets.

Challenges: As expected, storing terabytes of geographic information and searching on top of them is a cumbersome task. We need to incorporate summarization techniques such as multi-dimensional histograms to store the data as tight as possible using approximations while retaining as much of the information as possible. Building summarization techniques for geographic information with both spatial and keyword components is very challenging and requires the development of new techniques.

Products

1. Adaptive Event Detection with Time–Varying Poisson Processes Alexander Ihler; Jon Hutchins; Padhraic Smyth; none, 2006-06-15.
2. Counting Crowded Moving Objects vincent rabaud; Serge Belongie; 18th IEEE International Conference on Computer Vision and Pattern Recognition, 2006-06.

3. Domain-independent data cleaning via analysis of entity-relationship graph Dmitri Kalashnikov; Sharad Mehrotra; none, 2006-06.
4. Interactive Refinement of Filtering Queries on Streaming Intelligence Data Yiming Ma; Dawit Seid; Sharad Mehrotra; 2006 ISI, 2006-05.
5. Multi-perspective Video Analysis of Persons and Vehicles for Enhanced Situational Awareness Sangho Park; Mohan Trivedi; IEEE International Conference on Intelligence and Security Informatics (ISI-2006), 2006-05.
6. RAF: An Activation Framework for Refining Similarity Queries Using Learning Techniques Yiming Ma; Sharad Mehrotra; Dawit Yimam Seid; Qi Zhong; 11th International Conference on Database Systems for Advanced Applications (DASFAA 2006), 2006-04.
7. E – A Generic Event Model for Event-Centric Multimedia Data Management in eChronicle Applications Utz Westermann; Ramesh Jain; In Proc. of the IEEE Workshop on Electronic Chronicles (eChronicle 2006), Atlanta, Gerogia, 2006-04.
8. SAT: Spatial Awareness from Textual Input
9. Dmitri Kalashnikov; Yiming Ma; Sharad Mehrotra; Ramaswamy Hariharan; Nalini Venkatasubramanian; EDBT (Demo), 2006.
10. Events in Multimedia Electronic Chronicles (eChronicles). Utz Westermann; Ramesh Jain; Int'l Journal on Semantic Web & Information Systems, 2(2), 2006.
11. Towards a Common Event Model for Multimedia Utz Westermann; Ramesh Jain; IEEE MultiMedia, 2006.
12. Event Discovery in Multimedia Reconnaissance Data Using Spatio-Temporal Clustering. Bo Gong; Utz Westermann; Srikanth Agaram; Ramesh Jain; AAAI Workshop on Event Extraction and Synthesis, 2006.
13. Loss estimation online using INLET (Inter-based Loss Estimation Tool). Charles Huyck; Hung-Chi Chung; Sungbin Cho; Michael Mio; Shubharoop Ghosh; Ronald Eguchi; Sharad Mehrotra; 8th U.S. National Conference on Earthquake Engineering, San Francisco, 2006.
14. Suggestions for Effective Use of Remote Sensing Data in Emergency Response. Charles Huyck; Beverley Adams; Shubharoop Ghosh; Ronald Eguchi; 8th U.S. National Conference on Earthquake Engineering, San Francisco, 2006.
15. Remote Sensing Technology for Post-Earthquake Damage Assement. Beverly Adams; Charles Huyck; L. Gusella; C Wabnitz; S. Ghosh; Ron Eguchi; 8th National Conference on Earthquake Engineering, 2006.
16. Centralized web-based loss estimation tool. Charles Huyck; H. Chung; S Cho; M. Mio; S. Ghosh; Ron Eguchi; Sharad Mehrotra; SPIE Conference, San Diego, CA, 2006.
17. Streamlining Post-Earthquake Data Collection and Damage Assessment for the 2003 Bam, Iram, Earthquake Using VIEWS (Visualizing Impacts of Earthquakes With Satelites). Beverley Adams; Babak Mansouri; Charles Huyck; Earthquake Spectra, Volume 21, No. S1, pages S213–S218, 2006.
18. Low-Complexity Source Coding using Gaussian Mixture Models, Lattice Vector Quantization and Recursive Coding with Application to Speech Spectrum Quantization. A.D. Subramaniam; B.D. Rao; W.R. Gardner; IEEE Transactions on Speech and Audio Processing, 2005.
19. Robust Feature Extraction for Continuous Speech Recognition using the MVDR Spectrum Estimation Method. Dharanipragada S.; U. Yapanel; B.D. Rao; Technical Report, 2005.
20. Quality-driven approximate methods for integrating GIS data. Ramaswamy Hariharan; Michal Shmueli-Scheuer; Chen Li; Sharad Mehrotra;other, 2005.

21. Distributed Interactive Video Arrays for Event Capture and Enhanced Situational Awareness. Mohan Trivedi; Tarak Gandhi; Kohsia Huang; IEEE Intelligent Systems, Special Issue on AI in Homeland Security, Volume 20, Issue 5, Sept. - Oct., 2005.
22. Towards Rapid Citywide Damage Mapping Using Neighborhood Edge Dissimilarities in Very High-Resolution Optical Satelite Imagery. Charles Huyck; Berverly Adams; Sungbin Cho; Hung-Chi Chung; Ronald Eguchi;Earthquake Spectra, Volume 21, No. S1, pages S255–S266, 2005.
23. Permutation Models for Relational Data. Carter Butts; Technical Report, 2005
24. Exact Bounds for Degree Centralization. Carter Butts; Technical Report, 2005.
25. Emergent Coordination in the World Trade Center Disaster. Carter Butts; M. Petrescu-Prahova; Technical Report, 2005.
26. Radio Communication Networks in the World Trade Center Disaster. Carter Butts; M. Petrescu-Prahova; Technical Report, 2005.
27. Multimodal Systems in the Management of Emergency Situations. J. Cristoforetti; Thesis, 2005.
28. The CAMAS Testbed. V. Balasubramanian; D. Massaguer; Qi Zhong; Sharad Mehrotra; Nalini Venkatasubramanian; Technical Report, 2005.
29. Spatial queries over (imprecise) event descriptions. Dmitri Kalashnikov; Yiming Ma; Ramaswamy Hariharan; Sharad Mehrotra; Technical Report, 2005.
30. GAAL: A Language and Algebra for Complex Analytical Queries over Large Attributed Graph Data. Dawit Seid; Sharad Mehrotra; Technical Report, 2005.
31. PSAP: Personalized Situation Awareness Portal for Large Scale Disasters. Yiming Ma; Ramaswamy Hariharan; Amnon Meyers; Yonghua Wu; M. Mio; C. Chemudugunta; Sharad Mehrotra; Nalini Venkatasubramanian; P. Smyth; Technical Report, 2005.
32. Modeling Transportation Routines using Hybrid Dynamic Mixed Networks. V. Gogate; R. Dechter; J. Marca; C. Rindt; Technical Report, 2005.
33. Approximate Inference Algorithms for Hybrid Bayesian Networks in Presence of Constraints V. Gogate; R. Dechter; Technical Report, 2005.
34. Sparse Solutions to Linear Inverse Problems with Multiple Measurement Vectors. S. Cotter; B. Rao; K. Engan; K. Kreutz-Delgado; Technical Report, 2005.
35. Streamling Post-Earthquake Data Collection and Damage Assessment for the 2003 Bam, Iram, Earthquake using VIEWS (Visual Impacts of Earthquakes With Satellites). B. Adams, B Mansouri, C. Huyck. Earthquake Spectra, 21, p. S213-S218, 2006.
36. A Centralized Web-based Loss Estimation and Transportation Modeling Platform for Disaster Response. C. Hung, C Huyck, S. Cho, M. Mio, R. Eguchi, M. Shimozuka, S. Mehrotra. Proceedings of the 9th International Conference on Structural Safety and Reliability, Rome Italy, 2005-2006.
37. Multiperspective Thermal IT and Video Arrays for 3D Body Tracking and Driver Activity Analysis. S. Cheng, S. Park., M Trivedi. IEEE International Workshop on Object Tracking and Classification in and Beyond the Visible Spectrum, 2005-2006.
38. Real-Time Stereo-Based Head Detection using Size, Shape and Disparity Constraints. S. J. Krotosky; S. Y. Cheng; M. M. Trivedi; IEEE Intelligent Vehicles Symposium, 2005-06
39. Performance Evaluation of a Vision Based Lane Tracker Designed for Driver Assistance Systems. J. McCall; M. M. Trivedi; IEEE Proc Intelligent Vehicles Symposium, 2005-06.
40. Lane Change Intent Analysis Using Robust Operators and Sparse Bayesian Learning. J. McCall; D. Wipf; M. M. Trivedi; B, Rao; IEEE International Workshop on Machine Vision

for Intelligent Vehicles in Conjunction with IEEE International Conference on Computer Vision and Pattern Recognition, 2005-06.

41. Driver Activity Analysis for Intelligent Vehicles: Issues and Development Framework. S. Park; M. M. Trivedi; IEEE Intelligent Vehicles, 2005-06.

Contributions

1. Ontario Emergency Management Information Portal (beta site)
   http://www.disasterportal.org/ontario

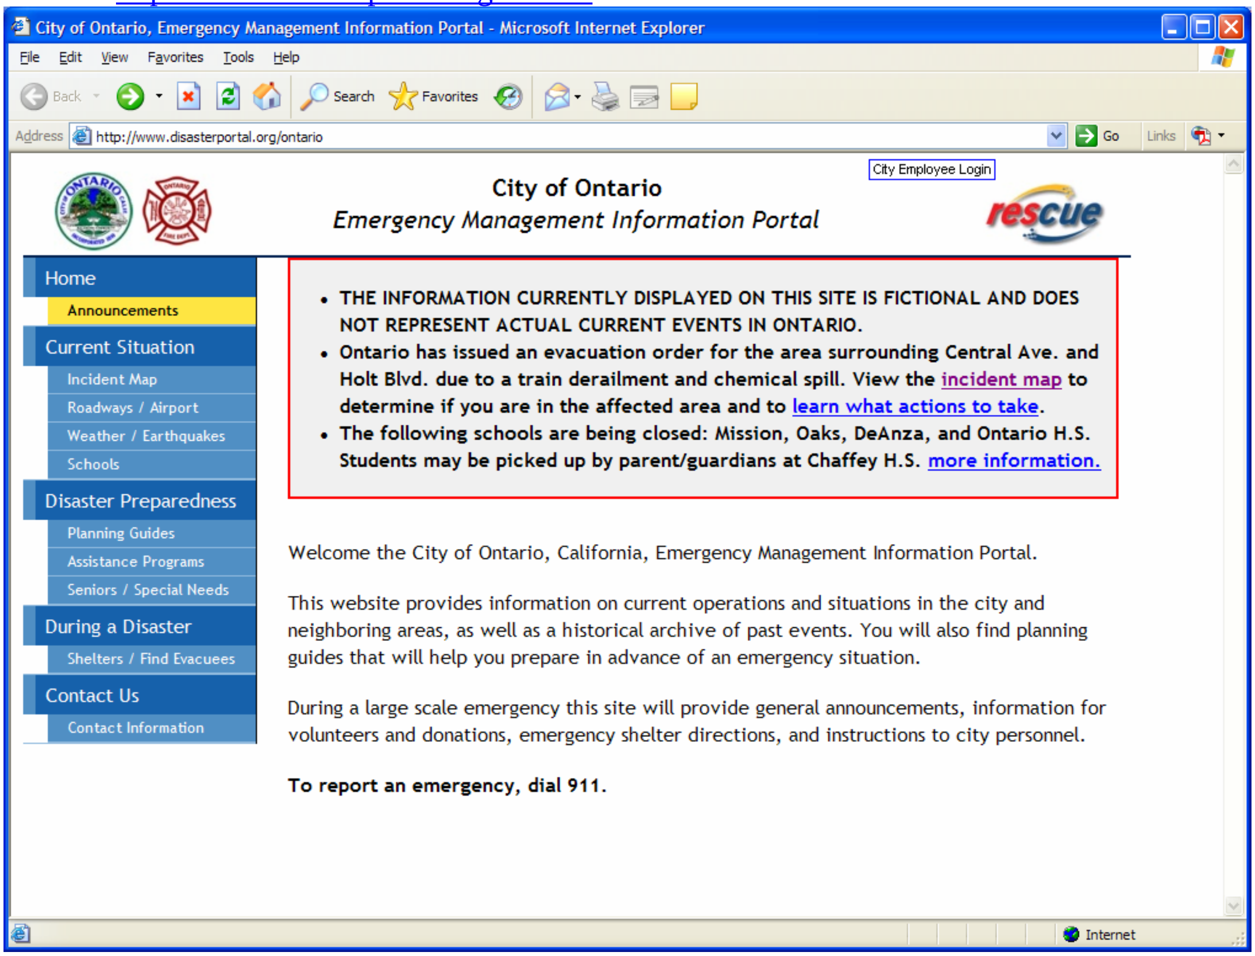

Ontario Portal

2. Responsphere Dataset Collection: http://rescue-ibm.calit2.uci.edu/datasets
3. Hurricane Portal

We have recently initiated the development of a national-scale Hurricane Portal in collaboration with Convera Inc. and ImageCat Inc. Our goal is to develop a portal for disasters (such as hurricanes) for real-world and large-scale use by exploiting the industrial strength search engine development technology provided by Convera Inc. We aim to provide 'faceted' navigation and search capabilities to analysts and post disaster information gathered from online information sources.

4. EvacPack

This artifact is a system for reconnaissance information capture. Our goal is to develop and demonstrate an end-to-end system for multi-modal situational information capture, with the eventual goal of the multi-modal data captured being useful for real-time situational awareness analysis or for post situation analysis and assessment.

For the EvacPack, we have completed the development of a mobile, wearable hardware sensor platform, with the following features and capabilities: audio and video information capture, GPS and other positioning indicator capture, heading and acceleration sensors, and video goggles, which will aid in visualization for a mobile user wearing the system. We are also currently developing a software infrastructure to manage multiple mobile sensor sources. This infrastructure includes: streaming (text data, audio, and video), data storage / management, temporal synchronization of various data, and visualization for analysts or people working in an EOC.

EvacPack

Policy-driven Information Sharing Architecture (PISA)

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 will 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 the past year we completed the construction of the derailment scenario that is the focus point for PISA. We have identified the RESCUE research that is relevant to the scenario. With the scenario in place, we began to identify the requirements for PISA in the context of the derailment. We also prepared for the sociology focus groups, planned for August 2006.

The most significant change from our initial plans is that the US government's Disaster Management Information System (DMIS) may serve as the underlying infrastructure for PISA, with policy management facilities layered atop DMIS. It is possible that DMIS may play the role we intended for ESB; we will be examining this point during the summer. Once this question is resolved, we can move on to integrating ESB and DMIS, if appropriate. It is clear now that SAMI will play a relatively minor role in the scenario, so we anticipate delaying the integration of SAMI with DMIS/ESB and instead concentrating on the addition of flexible authorization facilities to DMIS/ESB, as described below.

The information sharing policies that we have identified for the derailment scenario are relatively straightforward, i.e., they can be expressed in any reasonable policy language. We are investigating the possibility of layering them on top of DMIS. We have also determined that DMIS has rigid security requirements that may limit its effectiveness in a crisis. We plan to work to integrate RESCUE research on flexible authorization systems into DMIS over the coming year, and on trustworthy audit logs.

During the past year we have developed a number of approaches to enhance the availability of trust management infrastructure, even under attack and flash crowd conditions, and published those results in SATMAT06. We will be implementing those approaches in the coming year.

We continued our work on hidden credentials, which allow a service provider to send an encrypted message to a user in such a way that the user can only access the information with the proper credentials. Similarly, users can encrypt sensitive information disclosed to a service provider in the request for service. Policy concealment is accomplished through a secret splitting scheme that only leaks the parts of the policy that are satisfied. Hidden credentials may have relevance in crises involving ultra sensitive resources. They may also be able to play a role in situations where organizations are extremely reluctant to open up their systems to outsiders, especially when the information can be abused before an emergency even occurs. We have observed on the UCI campus that some buildings have lock boxes that are available to emergency personnel during a crisis. The management of physical keys is a significant problem. Hidden credentials have the potential to support digital lockboxes that store critical data to be used in a crisis. The private key used to access this information during a crisis may never have to be issued until the crisis occurs, limiting the risk of unauthorized access until the crisis occurs.nsert PISA Project verbiage.

Products

1. Traust: A Trust Negotiation-Based Authorization Service for Open Systems Adam Lee; Marianne Winslett; Jim Basney; Von Welch; The Eleventh ACM Symposium on Access Control Models and Technologies (SACMAT 2006), 2006-06
2. Synergy: A Policy-driven, Trust-aware Information Dissemination Framework. Ragib Hasan; Marianne Winslett; IEEE Intelligence and Security Informatics Conference 2006, 2006-05
3. Traust: A Trust Negotiation Based Authorization Service Adam Lee; Marianne Winslett; Jim Basney; Von Welch; The Fourth International Conference on Trust Management (iTrust 2006), 2006-05
4. Virtual Fingerprinting as a Foundation for Reputation in Open Systems Adam Lee; Marianne Winslett; The Fourth International Conference on Trust Management (iTrust 2006), 2006-05
5. Open Problems for Usable and Secure Open Systems. Adam Lee; Marianne Winslett; ACM CHI 2006 Workshop on Usability Research Challenges for Cyberinfrastructure and Tools, 2006-04.
6. Synergy: A Trust-aware, Policy-driven Information Dissemination Framework. Ragib Hasan; Thesis, 2005-12.
7. Adaptive Trust Negotiation and Access Control. T. Ryutov; L. Zhou; C. Neuman; T. Leithead; K. Seamons; 10th ACM Symposium on Access Control Models and Technologies, 2005.
8. Short Paper: Thor - The Hybrid Online Repository. T. van der Horst; K. Seamons; First IEEE International Conference on Security and privacy for Emerging Areas in Communications Networks, 2005.
9. Adaptive Trust Negotiation and Access Control for Grids. T. Ryutov; L. Zhou; C. Neuman; N. Foukia; T. Leithead; K. Seamons; 6th IEEE/ACM International Workshop on Grid Computing, 2005.

Contributions

TrustBuilder2

We continued our research on flexible authorization systems, which will serve as the building blocks for the flexible authorization facilities that we plan to demonstrate in the context of DMIS/ESB (and resulting in nine conference publications over the past year). To this end, we have been re-architecting TrustBuilder, our runtime system for authorization in open systems. TrustBuilder 2 will build on our insights obtained from using TrustBuilder over the past several years, by redesigning it to be more flexible, modular, extensible, tunable, and robust against attack. The design of TrustBuilder 2 is complete, and we are in the midst of implementation.

The hurdles hindering PKI deployment are also a huge obstacle to the deployment of some trust management solutions. We have begun exploring more lightweight mechanisms for establishing trust across security domains. Many crisis response organizations have limited information technology resources and training, especially in small to mid-size cities. We are giving more consideration to practical approaches for these environments.

Thor

We have developed Thor, a hybrid repository for storing and managing digital credentials, trusted root keys, passwords, and policies that is suitable for mobile environments. A user can download the security information that a device needs to perform sensitive transactions. The goals are ease of use and robustness. Our long-term goal is an architecture that emergency personnel will find easy to use to securely access sensitive data during a crisis.

We have completed the first integration of trust negotiation with a DBMS. The database system is aligned with a proxy that authenticates strangers outside the security domain according to rules and roles defined in the database system. This is a first step toward an information sharing architecture where organizations can use policy-based mechanisms to specify and control who has access to what resources.

LogCrypt

We have developed LogCrypt, which supports tamper evident log files using hash chaining. This system provides a service similar to TripWire, except that it is targeted for log files that are being modified. Often, an attacker breaks into a system and deletes the evidence of the break-in from an audit logs. The goal of LogCrypt is to make it possible to detect an unauthorized deletion or modification to a log file. Previous systems supporting this feature have incorporated symmetric encryption and an HMAC. LogCrypt also supports a public key variant that allows anyone to verify the log file. This means that the verifier does not need to be trusted. For the public key variant, if the original private key used to create the file is deleted, then it is impossible for anyone, even system administrators, to go back and modify the contents of the log file without being detected. During this past year, we completed experiments to measure the relative performance of available public key algorithms to demonstrate that a public key variant is practical. This variant has particular relevance in circumstances where the public trusts government authorities to behave correctly, and also benefits authorities by giving them a stronger basis for defending against claims of misbehavior. This technology may have relevance to more secure auditing during a crisis.

Nym

We are currently developing Nym, an extremely simple way to allow pseudonymous access to Internet services via anonymizing networks like Tor, without losing the ability to limit vandalism using popular techniques such as blocking owners of offending IP or email addresses. Nym uses a very straightforward application of blind signatures to create a pseudonymity system with extremely low barriers to adoption. Clients use an entirely browser-based application to pseudonymously obtain a blinded token which can be anonymously exchanged for an ordinary TLS client certificate. We designed and implemented a Javascript application and the necessary patch to use client certificates in the popular web application MediaWiki, which powers the popular free encyclopedia Wikipedia. Thus, Nym is a complete solution, able to be deployed with a bare minimum of time and infrastructure support. Nym currently authenticates clients based on their IP address. As part of a companion NSF project, we are beginning to explore how to leverage email authentication as a lightweight mechanism to authenticate and easily share information outside the local security domain. We should have results by the end of this year of the project, and it may provide a useful alternative for easy, policy-based sharing across organizations.

Customized Dissemination in the Large

We introduced a new form of dissemination that arises in mission critical applications such as crisis response, called flash dissemination. This year, we planned to develop algorithms/protocols for customized and flash dissemination in the Responsphere Smart-Space with heterogeneous devices, networks, users and information. We also expected to make progress on development of the RAPID flash dissemination system and test it under real-world conditions. We also planned to understand the problem of customized dissemination more formally, develop solutions to the key problems and integrate them into a usable implementation framework. We also expanded our work in several other directions and introduced new efforts that represent a more holistic approach to public dissemination as outlined in the strategic plan. This includes developing techniques that allow for customized and personalized dissemination information to the public at large (to individuals and organizations) through a peer-based publish-subscribe framework. We have also initiated a project on scalable and reliable information dissemination using heterogeneous communication networks (telephony, cellular, WiFi, mobile ad-hoc etc.) within the Responsphere testbed in addition to our effort in delivering information over traditional Internet based technologies. Two efforts in the social science arena focus on understanding information diffusion in interpersonal networks and across emergent multiorganizational networks. An understanding of such network structures and information flow through people in these networks guides the development and deployment of IT techniques for customized dissemination in wired/wireless networks. We describe in more detail the specific projects below.

Activities and Findings

Flash Dissemination in Heterogeneous Network

The goal of Flash Dissemination is rapid distribution (over the Internet) of varying amounts of data to a large number of recipients in as short a time period as possible. Given the unpredictability in its need, the unstable nature of networks/systems when it is needed and the medium data size of information, flash-dissemination has different concerns and constraints as compared to traditional broadcast or content delivery systems. First, it must be able to disseminate as fast (or faster) as current highly optimized content delivery systems under normal circumstances. But it should also be highly fault-resilient under unstable conditions and adapt rapidly in a constantly changing environment. In this study, we have developed a minimum-state gossip-based protocol called CREW. We implement the protocol in using a framework that we call RapID that provides key building blocks for P2P content delivery systems. Extensive experimental results over an emulated Internet testbed show that CREW is faster than current state-of-art dissemination systems (such as BitTorrent, Splitstream and Bullet), while maintaining the stateless fault-tolerance properties of traditional gossip protocols. Further, we have also investigated centralized, global-knowledge protocols for dissemination. These protocols are much faster than localized randomized protocols and may be useful in situations where network and systems are stable and dissemination may be repeated multiple times to the same set of recipients, thus amortizing the knowledge collection and dissemination plan generation costs. Our key findings are as follows:

- We have tested various dissemination systems and protocols on our Internet testbed. The systems include BitTorrent, Bullet, Splitstream and lpbcast. The first three systems are highly optimized systems for large content delivery in heterogeneous networks. Lpbcast is a gossip-based protocol more geared towards fault-tolerance.

- In comparison to these systems, CREW, our flash-dissemination protocol achieved much faster dissemination, under various heterogeneity settings such as heterogeneous latencies, bandwidths and packet loss rate.

- CREW is a gossip-based protocol. A major implication of this is that gossip-based protocols can be so designed so that they can achieve (and even better) the performance of optimized dissemination systems without compromising on their fault tolerance properties.

- Gossip-based protocols are inherently simpler to design and more fault-tolerant. Thus, CREW has dual advantage of lower maintenance of system code and higher resilience during crisis scenarios while still being able to disseminate data very quickly.

- Two key factors are needed to make gossip-based protocols achieve fast dissemination. First is the reduction in data overhead and second is high concurrency.

- Low data overhead can be achieved using two techniques. First meta-data needs to be used so that nodes only get information that they are missing. A more novel finding is the use of random walks on overlays to achieve a near real-time constant overhead membership service.

- High concurrency can be achieved by making all nodes active as soon as possible (high inter-node concurrency) and by using a high-performance and scalable middleware layer to achieve high intra-node concurrency.

- High concurrency leads to congestion in the network slowing down dissemination. Thus high concurrency needs to be autonomically adaptive. We implemented a congestion recognizer and backoff mechanism in CREW using theory from random sampling to achieve this. Experimental results show that this combination of high concurrency coupled with intelligent backoff results in CREW’s superior perforamce.

- We have also come up with centralized dissemination heuristics that can outperform randomized approaches when global knowledge is available. These heuristics can be employed in situations where dissemination needs are known beforehand.

- We have started investigating how overlay properties affect dissemination speeds. Preliminary investigation shows that highly skewed overlays can significantly impact the dissemination speed. Making dissemination systems agnostic to overlay properties is an important step for several reasons. It opens up new possibilities in design of overlay construction which has implications for constructing fault-tolerant overlays. Also, it has implications in the use of dissemination protocols in other systems where the overlay is beyond the control of the dissemination protocol.

- We will also investigate how the dissemination protocols behave in highly dynamic environments where a large number of nodes can join and leave simultaneously. This has implications ranging from building a ‘catastrophe tolerant’ dissemination system to building a P2P web-server for handling large unpredictable flash crowds.

Flash Dissemination in Heterogeneous Networks

Existing pub/sub systems usually consists of a set of pub/sub servers (brokers) connected to each other through an overlay network. Each broker acts as a local access point and manages the set of clients connected to it. Publishers and subscribers connect to one of the brokers to send or receive publications and subscriptions. These systems construct a spanning tree on the pub/sub overlay network to forward subscriptions and publications. This allows avoiding diffusion of content in parts of the pub/sub network where there are no subscribers and prevent multiple deliveries of events. The other reason for using spanning tree is to exploit subscription aggregation to reduce subscription maintenance overhead.

The main goal of the most of the existing approaches in publish/subscribe systems is to increase scalability of the pub/sub system through reducing each broker’s load which consists of subscription storage and content matching. These systems aim to reduce the total cost of matching and routing events and can scale to large number of subscribers. However, in many pub/sub applications notification dissemination time, the time between publishing a notification and delivering it to all of the interested subscribers, also plays a critical role. We are using content-based pub/sub as a communication infrastructure for a customized alert dissemination system. In this application timely dissemination of notifications to the interested receivers is the primary goal. While existing content-based pub/sub systems can scale well and disseminate notifications in a reasonable time and communication cost which is sufficient for most of applications, we believe our specific application and similar systems need faster notification dissemination. Therefore, we look at content-based pub/sub system from a different angle. The main goal of the pub/sub system in this view is to disseminate notifications among interested subscribers as fast as possible. To achieve this we need to address the factors that involve in the notification dissemination process and exploit all the available resources. The main operations that affect the performance of a content pub/sub are content matching and content routing.

- We propose a new approach to content based pub/sub that disseminates messages from publishers to subscribers along a shortest path. By using the shortest path, our approach is able to significantly reduce the delivery time of notifications under diverse load. Besides improving dissemination speed, our approach naturally exploits the inherent parallelism of the overlay network and reduces the likelihood of links becoming bottlenecks by distributing notification dissemination task to a larger number links compared to the existing approaches.
- We develop two main approaches both of which use shortest path for message delivery. These approaches differ in the techniques used to prune the content dissemination tree to prevent forwarding published content towards part of the overlay network where there are no subscribers. Our first approach, LACRA, aims to reduce subscription maintenance by aggregating subscriptions in every broker based on the link that the subscriptions are received from. Based on this approach, we present two algorithms to disseminate subscriptions and content. The second approach, BACRA, we propose aims to accurately route content by having subscriptions

broadcast to all brokers. We propose two different content routing algorithms based on our second approach.

• We also propose a novel approach for content representation which exploits subscription covering and merging and speeds up the matching operation. This is done by applying multidimensional indexing techniques in pub/sub system.

• Through comprehensive simulation we show that our proposed approach significantly speeds up the content dissemination process while the dissemination traffic is less or at most the same as the existing systems. Also our simulations show that content representation technique that we propose can achieve considerable reduction in subscription maintenance and content matching load.

Information Dissemination in a Heterogeneous Wireless Environment

Wireless networks (e.g., cellular, Wi-Fi) extend wireline networks in warning and notifying large number of people (both the public and first responders) in crisis situations. People with handheld devices (e.g., cell phones, PDAs) not only receive emergency alerts, but also share warnings and other related information between each other via ad hoc networks. Information dissemination in this context needs to reach maximum number of intended recipients (e.g., affected people) within the shortest possible time; and the data to be disseminated can be quite large (e.g., image, voice, etc.). In this work, we study fast, reliable, and efficient dissemination of application-generated data in heterogeneous wireless networks. We consider coverage (the percentage of intended recipients that receive the information), redundancy (the time it takes for people to receive the information) and energy consumption (on handheld devices) as the primary metrics. We develop efficient dissemination strategies that are not only fast and reliable, but also resilient to network congestion and recipients' mobility. We propose protocols that manage the dissemination of data with large size. We also investigate exploiting multiple radio interfaces, hybrid networks as well as mobility for faster dissemination.

Key findings include:

1. We have obtained a more concrete understanding of the distinct needs of application-level data dissemination in wireless networks, i.e., high coverage (reaching the maximum number of recipients), and low latency (within shortest possible time), etc..
2. We have reexamined the network stack for application data broadcast, and evaluate the choices of primitives on each layer, for instance, MAC-broadcast Vs MAC-unicast at the MAC layer, empty routing Vs on-demand routing at the network layer, etc..
3. Through experimental studies, we have shown the performance tradeoffs between using MAC-broadcast and MAC-unicast as the primitives for wireless data broadcast. Specifically, MAC-broadcast-based dissemination approaches have better performance in terms of high coverage and short latency, whereas MAC-unicast-based dissemination approaches save energy.
4. We have pinpointed the advantages and drawbacks of existing dissemination schemes, and have reexamined the well-known broadcast storm problem from the perspective of application-data dissemination. We have found that blind flooding is an effective protocol, which achieves both high coverage and low latency.
5. We have briefly studied the power properties of wireless interfaces in MAC-broadcast and MAC-unicast transmissions, including how they consume energy in transmitting, receiving, idling as well as in reduced-power state, etc..

6. We have shown the good performance of single-spanning-tree-based dissemination schemes, and the nice property of a center-rooted spanning tree. We have developed a distributed center-finding algorithm, which quickly finds the approximate central node of a network using small amount of messages.
7. We have proposed to use random unicast-based background protocols to save battery on participating handheld devices during disseminations. Unicast transmissions can put neighboring nodes into reduced-power state, in which less energy is consumed as compared to usual.
8. We have proposed to exploit cellular/Wi-Fi combined networks to achieve quick, scalable and location-aware critical-information dissemination, using new technologies such as cell broadcasting and Wi-Fi cell phones. We have shown the feasibility and advantages via a simulation-based demo.

Achieving Communication Efficiency through Push-Pull Partitioning of Semantic Spaces to

Disseminate Dynamic Information

Many database applications that need to disseminate dynamic information from a server to various clients can suffer from heavy communication costs. Data caching at a client can help mitigate these costs, particularly when individual push-pull decisions are made for the different semantic regions in the data space. The server is responsible for notifying the client about updates in the push regions. The client needs to contact the server for queries that ask for data in the pull regions. We call the idea of partitioning the data space into push-pull regions to minimize communication cost data gerrymandering. In this study we present solutions to technical challenges in adopting this simple but powerful idea. We give a provably optimal-cost dynamic programming algorithm for gerrymandering on a single query attribute. We propose a family of efficient heuristics for gerrymandering on multiple query attributes. We handle the dynamic case in which the workloads of queries and updates evolve over time. We validate our methods through extensive experiments on real and synthetic data sets.

A Diffusion of Crisis Information through Interpersonal Networks

The goal of this effort is to understand the information diffusion process on hypothetical population of persons within a region. In this preliminary work, the network is loosely typical of what one might expect from telephone contacts; vertices are scaled by the expected completeness of the information they would be expected to receive from a serially transmission process (e.g., word of mouth) in which each concept has a chance of being lost during each iteration. Factors studied include the number of steps from the point of origin and the time to first contact the node. Interesting observations include:
1. Spatial character of the process: the movement of information seems to be much more uneven than might be expected from a purely spatial model: information does occasionally "tunnel" to spatially remote parts of the network, where spatially local clusters of informed actors often emerge. Within regions which are generally well-informed, one generally finds a few "stragglers" who are notified very late in the diffusion history. One expects that this effect would be amplified by additional sources of heterogeneity (e.g., age, race, language, etc.), but it emerges even without those complicating factors.

2. Path lengths: Information often takes a fairly circuitous route in reaching its destination. Even where network diameters are on the order of 3-4, realized path lengths of 12-14 (or longer) may be observed. This has important implications for the nature and quality of information which is received by individuals in the network; in particular, few complex signals can persist for that number of steps.
3. Notification time: Information quality declines in first notification time, spatial distance from the origin, and social distance from the origin. Simple signals may persist over fairly long chains, but detailed information is likely to be concentrated within a reasonably narrow socio/spatio/temporal envelope around the originating individual.

Our preliminary efforts indicate that the behaviors are consistent with the field literature on initial propagation of information during crises of various sorts, which is encouraging. Two possible applications of this sort of work is in:
1. Reverse 911 systems for mass dissemination where one could exploit the natural process of information diffusion by strategically placing reverse 911 calls (or other alerts) so as to maximize the diffusion of accurate information within a target population.
2. Modeling phenomena such as mass convergence, auto-evacuation which result from the decisions of persons who are alerted to an impact event in their vicinity by projecting which groups are most likely to engage in such behavior within the immediate post-impact period. This in turn can be used to predict resource needs (communication, transportation etc.) and plan the deployment of resources.

Understanding Emergent Multiorganizational Networks in Crisis

In the wake of disasters, previous researchers have identified emergent systems among groups and organizations. When existing organizational networks are unable to cope with an event, emergent multiorganizational networks (EMONs) develop to meet the needs of the affected community. Using data collected from the Sept 11th WTC attacks, this effort analyzes the multiorganizational networks that formed after the WTC attacks and develops a methodology for extracting EMONS that can be applied in practice.

Such analyses apply to dissemination of information to the public since they help learn which organizations provide greater numbers of linkages between like organizations and dissimilar organizations. This may help to identify which organizations have the greatest ability to distribute information to others due to their role and place in the network structure. For instance, there may be certain non-profits or non governmental organizations that serve to link local organizations with diverse constituencies that otherwise would not be reached through more traditional channels. By identifying these central organizations, one can be more certain that information is being sent to the peripheral local organizations and thereby reaching the most vulnerable populations such as those served by small, local nonprofits and human service organizations.

Following the September 11 attacks on the World Trade Center, data was collected from situation reports, status updates, newspaper articles, news transcripts and field notes to document emergent multiorganizational networks (EMONs) which developed in response to the WTC disaster. These field documents and newspaper articles were coded and recorded to represent the more than 6,600 records of interaction between organizations involved in the initial 12 days following the event. The resulting data set represents one of the largest efforts ever undertaken to capture EMONs in the immediate post- disaster environment. Due to phenomena such as mass convergence to the impact site, the need for improvisation in a changing environment, and the presence of conditions that exceed the capabilities of the pre-existing response system, EMONs typically emerge to coordinate response activities among the many organizations involved in the response process. Relatively little is known, however, regarding the structure of such networks, or the determinants of interaction with them.

In our work, we examine the probability of interaction between organizations based upon attribute data, including type (i.e., government, non-profit, profit, collective), scale of operation, and degree of autonomy. using exponential family models we estimate the extent to which organizations work with similar versus dissimilar alters (i.e. non-profits working with non-profits or government organizations working with for-profit organizations). In addition, we investigate the question of whether these effects differ depending upon the functional tasks in which the organizations are involved. These results shed light on the emergence of coordination among organizations of various types in this post-disaster environment of September 11.

Our work also presents a step-by-step methodology to measure and extract EMON's using qualitative data for quantitative network analysis. This methodology addresses issues related to data sources, coding, network extraction, and analysis. Issues related to data sources include common sources and their properties, sampling issues, and data management, including cataloging and archiving data. Data coding and network extraction will be concerned with the identification of vertices and their relationships, or edges. Finally, considerations for analysis will address measurement constraints and errors. The findings from these networks can provide more general insights for future large-scale disaster events.

Products.

1. Deshpande, M., Xing, B., Lazaridis, I., Hore, B. Venkatasubramanian, N., Mehrotra, S. *CREW: A Gossip-based Flash-Dissemination System*. ICSDCS (International Conference on Distributed Computing Systems), 2006.
2. Achieving Communication Efficiency through Push-Pull Partitioning of Semantic Spaces to Disseminate Dynamic Information Amitabha Bagchi; Amitabh Chaudhary; Michael Goodrich; Chen Li; Michal Shmueli-Scheuer; Technical Report, 2006-03.
3. Information Collection Services for QoS-aware Mobile Applications. Qi Han; Nalini Venkatasubramanian; IEEE Transactions on Mobile Computing, 2005.
4. Real-time Collection of Dynamic Data with Quality Guarantees. Qi Han; Nalini Venkatasubramanian; IEEE Tran. Parallel and Distributed Systems, in revision, 2005.
5. Customized Dissemination in the context of emergencies. A. Ghigi; Thesis, 2005.
6. Scalable, Flash Dissemination of Information in a Heterogeneous Network. Mayur Deshpande; Nalini Venkatasubramanian; Sharad Mehrotra; Technical Report, 2005.
7. Dissemination of Evacuation Information in Crisis. A. Ghigi; Technical Report, 2005.
8. I.T. Support for Information Dissemination During Crises. Mayur Deshpande; Nalini Venkatasubramanian; Sharad Mehrotra; Technical Report, 2005.

Contributions

1. RAPID: http://www.ics.uci.edu/~mayur/rapid.html
2. Real-Time Crisis Alert System: http://www.ics.uci.edu/~projects/dissemination/

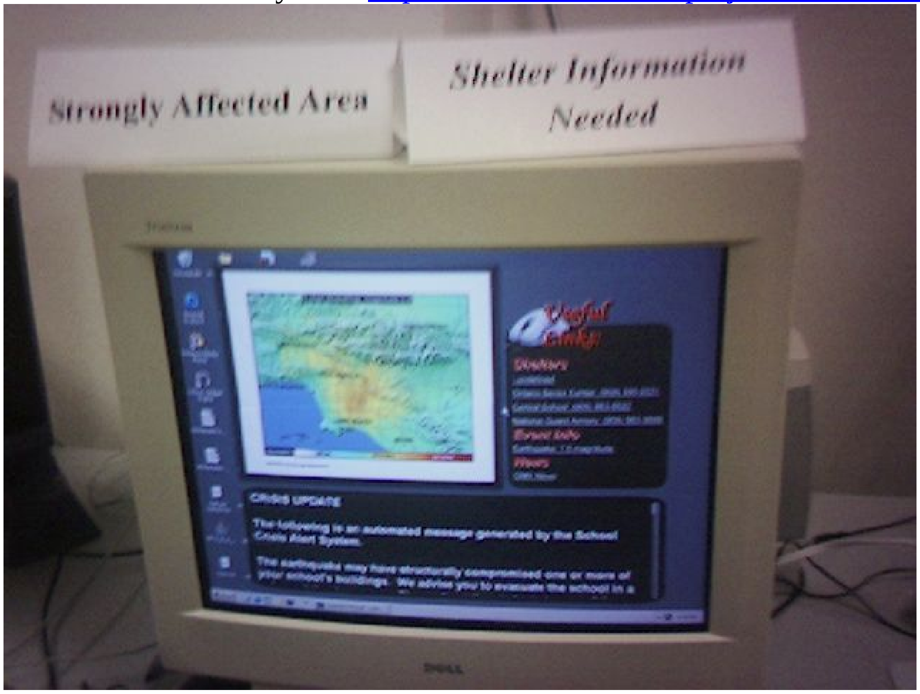

Real-Time Alert System

Privacy Implications of Technology

Privacy concerns in infusing technology into real-world processes and activities 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 make persistent 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 the society. Our objective is to understand privacy concerns in adopting technology from the social and cultural perspective, and design socio-technological solutions to alleviate such concerns.

Activities and Findings.

Understanding Privacy Concerns of Technology Infusion

A Study of Privacy and Utility for Emerging Surveillance Technologies: Surveillance technology automates human perception capabilities in a manner that alleviates the constraints of the physical world. In other words, it facilitates time-, space-, and cognition-shifting in ways that can both greatly enhance capabilities as well as being exploited to violate privacy rights and expectations. There are opposing arguments as to whether the benefits of such technologies outweigh the costs. Among these considerations, the real and potential privacy risks are paramount and often define the course and scope of surveillance technology deployments. Specifically, the battle is often fought between the privacy rights lobby and government/law enforcement and corporate entities. The RESCUE privacy team (in a recently established collaborative project with Erin Keneally, M.F.S., J.D.) launched a study that will inform the cost-benefit debate by assessing how surveillance technologies can be designed and deployed to minimize the potential privacy- invasive applications of surveillance technology, while simultaneously realizing it's utility.

Ethical considerations of business continuity and disaster recovery: We also explored the ethical foundations (e.g., philosophical and theoretical framework) and future considerations of individual privacy, data privacy, data security, and data custody within the domain of business continuity and disaster recovery.

Quantifying Privacy and Privacy-Preserving Mechanisms

Systematic Search for Optimal k-Anonymity: In this work, we studied the problem of publishing individual centric multi-attribute data for the purposes of data release. All explicitly identifying data attributes such as name, ssn etc. are removed from the released data. But the problem is that a combination of the remaining attributes can be unique and thereby act as an identifying attribute. For instance, the combination of date of birth, zip-code (of residence) and age (in years) is unique for any individual with a very high probability. Consider a released data set containing confidential medical data about individuals is to be published. Now, if the data set is released along with the date of birth, zip-code and age attributes without any form of modification, it is highly likely that individuals can be uniquely identified using these three “quasi identifier” attributes and thereby lead to privacy violation by disclosing the confidential medical information about the identified individual. The popular class of techniques that are used to modify such data is called “generalization” where the exact values of the quasi-identifiers are replaced by a more generalized value. For example, the exact age may be replaced by an age range or a categorical data value might be generalized to a higher level category if such a hierarchy (taxonomy) is available. In general, the available set of such generalization schemes is very large. Also each such modification scheme affects the “quality” of the modified data set. Quality of the modified data set is generally captured quantitatively by some cost measure which reflects its utility for some target data mining activity. For instance, if the goal of data mining is to train a classifier on the dataset in order to learn the association between the quasi identifying attributes and the medical information of individuals, an appropriate information-loss/cost metric would be one that reflects the quality (error-probability) of such a learnt model. In general the approach taken for generalization is to consider some family of generalization schemes and try to find the optimal one from this family which minimizes the information loss metric. If we visualize the data in a multidimensional space where each attribute corresponds to an unique dimension, then most of the generalization schemes proposed till date, can be considered to be space partitioning schemes where the data space is split at various attribute values along various dimensions and these splits are allowed to go all across the data space. That is, the partition boundaries go all the way through from one end of the data space to another. In contrast our partitioning scheme can be considered multidimensional which allows more flexible space partitioning schemes where such a constraint is not imposed on the partition boundaries. Our family of partitioning schemes is a superset of the families considered by single dimensional partitioning. The class of partitioning that we generate are also called “hierarchical” or “guillotine” partitioning schemes. More specifically, we achieved the following:

a. Developed a novel enumeration scheme for duplicate free enumeration of all hierarchical partitioning of the space for a given set of splits.

b. Developed generic pruning strategies for efficient search for optimal partitioning schemes that minimize a wide class of information loss metrics under variety of different constraints.

c. Carried out extensive experimentation on both real-life and synthetic datasets and empirically shown how our technique improves upon all previous approaches and is applicable for a wide variety of metrics.

We show how our approach is quite general and works for a large class of cost functions and constraints. We use the priority queue data structure to trade-off space for efficiency (i.e., find good solutions faster). In fact, the priority queue data structure can be used to run the algorithm in two modes, one in which the goal is to search for the optimal solution and the other in which the goal is to search for a t-approximate solution where t is pre-specified. The two modes lead to different execution patterns in general. We also propose a couple of new generic cost/information-loss metrics and describe the sufficient properties of the class of cost measures and respective constraints that can be optimized using our search-tree based approach. We study two cost functions: the discrimination metric (DM) and the volume metric (VM) with a variety of constraints like the minium attribute range constraint (e.g., age should not be specified to less than a range of 10 years), the "minimum entropy of sensitive value" constraint (e.g., minimum diversity of sensitive values within an anonymized class is guaranteed) and the simple k-anonymity constraint. We compared the performance of our algorithm with that of greedy approach proposed in a related paper and show the advantage of our generic, flexible approach over theirs.

Location Privacy: Privacy has to be examined in the dynamic context where data is generated or maintained. Specifically, location data subjects to spatial constraints as well as other correlations that allow inference among data. Hence, for instance, when requesting a location based service, a pseudonymous user should consider not only other users requesting services at the same time, but also the requests (which contain location information) she sent out earlier.

Storing and Querying Data in Untrusted Environments: We extend our previous work which analyzed in depth, the bucketization approach for supporting single dimensional range queries. In the current extension we look at the case when we want to support multidimensional range queries and join queries in the DAS model. We are currently investigating the problem of optimal bucketization in multidimensional space to support such queries, and analyzing the threat of disclosure in such model and exploring new ways to trade-off disclosure risk versus performance. All of these issues are significantly more complex in the multidimensional model as compared to the single dimensional case and is providing us a much more generic perspective of the nature of the disclosure problem in the DAS model.

Privacy-Preserving Observation Systems

Privacy-Preserving Pervasive Systems (P3): Building on the work done in the PaDOC system (video surveillance), we derived insights into the nature of privacy preservation in pervasive environments. Specifically we achieved the following: Extended the basic model for specifying certain types of policies in pervasive spaces to support more complex policies that rely on historical user-data gathered in the pervasive space; Studied this from the point-of-view of specific scenarios on-site (smart office-space/coffee room); Proposed an event automaton-based execution model for user-specified policies; Proposed a system architecture and designed communication protocols for detection of events; Proposed a very specific notion of privacy in the context of the architecture – i.e. anonymity; Determined channels of inference in the proposed system that may lead to privacy violations; Identified the various constraints that need to be imposed on data representation and communication protocols to ensure individual’s privacy in such a system; Modeled the system design problem (privacy-performance) as a constrained optimization problem and proposed some heuristic solutions for the same; Experimentation on the proposed protocols under varying conditions (e.g. levels of anonymity, structure of composite events, etc.).

Privacy-Preserving Data Sharing Systems

Protecting Individual Information Against Inference Attacks in Data Publishing: In many data-publishing applications, the data owner needs to protect sensitive information pertaining to individuals, such as the disease of a patient. Meanwhile, certain information is required to be published. The sensitive information could be considered as leaked, if an adversary can infer the real value of a sensitive entry with a high confidence. There are various methods using which the adversary can infer sensitive values. In this paper we study how to protect data when an adversary can do inference attacks using association rules derived from the data. We formulate the inference attack model, and develop complexity results on computing a safe partial table. We classify the general problem into sub-cases based on the requirements of publishing information, and propose the corresponding algorithms for finding a safe partial table to publish. We have conducted an empirical study to evaluate these algorithms on real data.

pVault/Delegate: Explored architecture for achieving secure mobile access to personal data. In this architecture the users outsource their personal information to a remote service provider who is in charge of providing storage and data access services. The heterogeneous personal data of users is captured in the form of XML documents. The service provider itself is untrusted and therefore the data is encrypted before being outsourced. This architecture allows the user to access their data from any trusted computer connected to the internet. This architecture was implemented as a software artifact called Pvault. Pvault has been running successfully for over a year. We proposed a new architecture for accessing websites securely from untrusted machines (delegate). This architecture allows the users the functionality of accessing their websites from untrusted machines without revealing their secrets. This architecture specifically prevents key logging, shoulder snooping and password sniffing attacks. Currently, this architecture is under development.

Products.

1. Intelligence and Security Informatics (ISI) 2006. Prof. Sharad Mehrotra was the program Co-Chair for the ISI Confererence. ISI brings together researchers and practitioners from industry, academia, and government and the conference covered topics including homeland security, terrorism informatics, and crisis response
2. Cybertrust 2005. Prof. Sharad Mehrotra organized the NSF Principal Investigators Meeting for the Cybertrust Program. Cybertrust brings together researchers in security, database security, privacy, and cryptography together 2005.
3. Rescue Distinguished Lecture Series (2005-2006).
4. Rescue Next-Generation Search Series (2005-2006).
5. Davison, C. *Ethical Considerations of Business Continuity and Disaster Recovery Planning*. International Academy of Business and Technology (IABT) 2006. Conference.
6. Mykletun, E. & Tsudik, G. *On using Secure Hardware in Outsourced Databases*. International Workshop on Innovative Architecture for Future Generation High-Performance Processors and Systems, 2005.
7. Jammalamadaka, R., Mehrotra, S. & Venkatasubramanian, N. *PVault: A Client Server System Providing Mobile Access to Personal Data*. StorageSS 2005., Held in conjunction with ACM CCS 2005.
8. Hore, B., Wickramasuriya, J., Mehrotra, S., & Venkatasubramanian, N. *Privacy-Preserving Event Detection for Media Spaces.* Technical Report, 2006.
9. Wickramasuriya, J., Datt, M., Mehrotra, S., & Venkatasubramanian, N. *Secure Logging for Privacy-Preserving Data Collection in Pervasive Spaces,* Technical Report, 2005.

10. Hore, B., Wickramasuriya, J., Mehrotra, S., & Venkatasubramanian, N. *Pervasive, Not Invasive: Privacy Preservation for Pervasive Spaces*. Rescue Journal, 2005.
11. Jammalamadaka, R., Van der Horst, T., Mehrotra, S., Seamons, S., Venkatasubramanian, N. *Delegate:A Proxy Based Architecture for Secure Website Access from an Untrusted Machine.* Technical Report, 2006.
12. Hore, B., Jammalamadaka, R., Mehrotra, S. *Flexible Anonymization for Privacy Preserving Data Publishing: A Systematic Search based Approach.* Technical Report, 2006-03.
13. Dilmaghani , R., Manoj, B., Rao, R. *Emergency Communication Challenges and Privacy*. 2006 ISCRAM, 2006-05.
14. Li, C., Shirani-Mehr , H., Yang, X. *Protecting Individual Information Against Inference Attacks in Data Publishing.* Technical Report, 2006-03.
15. Holt, J. *LogCrypt: Forward Security in Public Verification for Secure Audit Logs*. Australasian Information Security Workshop, 2006.
16. Mykletun, E. Tsudik, G. On using Secure Hardware in Out-sources Databases. International Workshop on Innovative Architecture for Future Generation High-Performance Processors and Systems, 2005.
17. Jammalamadaka, R., Mehrotra, S. *Querying Encrypted XML Documents.* Technical Report, 2005.
18. Park, S., Trivdei, M. *A Track-based Human Movement Analysis and Privacy Protection System Adaptive to Environmental Contexts*. IEEE International Conference on Advanced Video and Signal based Surveillance, 09/2005.

Contributions.

1. PaDOC (A Framework for Privacy-Aware Data Collection) [http://p3.ics.uci.edu]
2. pVault (Secure Password Manager) [http://www.itr-rescue.org/pVault/]

Robust Networking and Information Collection

The Extreme Networking System (ENS) is the artifact from the robust networking and information collection project and the ENS system is a robust networking infrastructure which provides computing, communication, and intelligent information collection, management, and maintenance systems, for use on site at ground zero, in the event of emergencies. Given the possible environmental constraints such as lack of electric power, partial or full unavailability of fixed communication networks, and the presence of heterogeneous sets of communication technologies, designing an ENS is a challenging task yet ENS is critical to information collection, management, and dissemination process, which are essential components for creating situation awareness. Thus ENS is an enabler for the higher layer functions that support situational awareness that are being *d*eveloped as part of this project.

The primary goal behind the artifact is to provide an efficient, reliable, and scalable network infrastructure for aiding the emergency response activities when the existing networking infrastructure fails or becoming non-functional. Other objectives of this artifact it to make the application of research ideas developed as part of the robust networking and information collection research project in RESCUE.

The ENS artifact is a three-level hierarchical network architecture has been developed and deployed as part of this project. The first level is formed by the user or responder devices which, to accommodate the needs of first responders, can be quite heterogeneous. The second level is formed by a wireless mesh network plane which can provide high reliability and fault tolerance. The third level is formed by a variety of multiple long haul backbone networks such as cellular and satellite networks. The gateway nodes act as the bridge between the wireless mesh planes to the backbone networks. In addition to the three levels of networking modules, ENS bundles a set of application layer solutions for information collection, management, and intelligent dissemination. A portable ENS node, CalMesh node, is the major component of ENS. A CalMesh node can incorporate multiple technologies and interfaces to support the other two hierarchies in addition to performing its primary task at the wireless mesh network plane. Each CalMesh node has the capability to provide additional information such as geo-location information which helps in generating situational awareness and contextual information. The ENS also provides localized and customized information management and maintenance resources such as localized web services at ground zero. ENS has inbuilt capability to provide adaptive content processing and information dissemination to the first responders and the victim population. The internal view of CalMesh node is shown in Figure (b), and Figure (c) depicts the portability of CalMesh node. The current version of the ENS architecture has been used under several trial experiments.

The ENS architecture differs from other network architectures as it utilizes several advanced features such as Always Best Connected (ABC) paradigm, bandwidth aggregation techniques, load balancing mechanisms, and localized web-based information collection, management, maintenance, and intelligent dissemination system, besides being an example of next generation hybrid wireless network architecture.

Activities and Findings.

Powerline Networking

The networking researchers at UCI instrumented the CalIT2 building with sensors that communicate over the existing powerlines (i.e., copper-based power grid system). The goal is to provide low bandwidth sensor information over yet another layer 1 infrastructure. During a disaster event, this could provide a redundant source of information outside the normal layer 1 communications infrastructure.

The powerline system is provided by Echelon Corporation. The system provides approximately 115Kbps bandwidth under most conditions. An interesting finding is that during power failures, the bandwidth actually increases due to less electrical noise on the copper wiring.

The researchers are experimenting with temperature and humidity sensors on the powerline network. These sensors log to a web addressable platform that can be queried for real-time sensor information. Additionally, the sensors log to a DB2 database so that all past sensor readings may be queried.

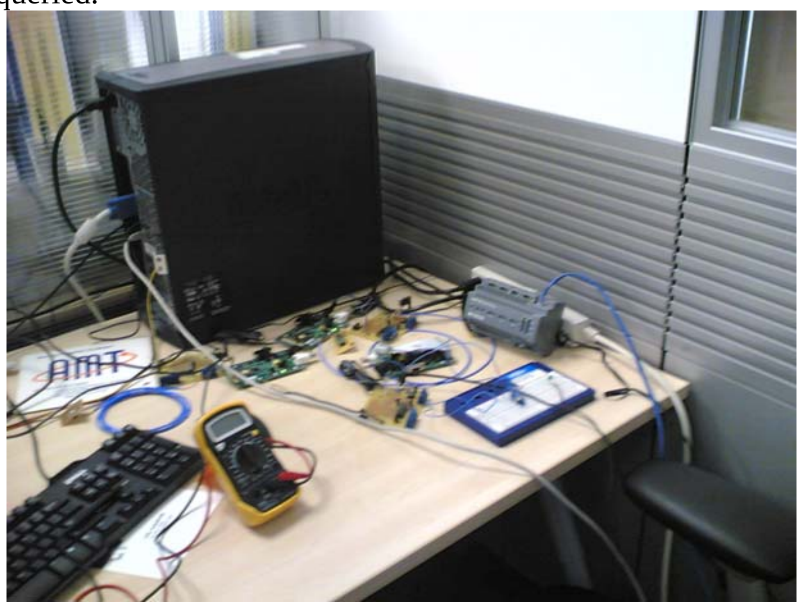

Powerline Networking System

GasLamp Quarter Wireless Mesh Deployment

The Gaslamp Quarter is an approximately three-by-eight square block area from 4th Street to 6th Street (west to east) and E Street to L Street (north to south) in downtown San Diego, CA. As part of the Mardi Gras exercise on February 28, 2006, both the SDPD command post and the UCSD Tech Ops Center were located inside the event perimeter at a hotel near the southern end of the covered area.

The Tropos network system consisted of five Tropos 5110 outdoor routers, five outdoor remote controllable on/off switches and five Tropos power cords; these were installed five days in advance of Mardi Gras (and still remain in the testbed). The system included one gateway node which was connected to a 3 Mbps wireless backhaul. This gateway node was installed on the roof of a two-story business building. The other four Tropos nodes were installed on street-corner lamp posts. CalMesh networking nodes were used to extend the coverage of the wireless network: Three CalMesh nodes were deployed to connect the command center and southern-most camera to the satellite internet connection.

Tropos Node affixed to a lamppost in the GLQ

VPN clients, built by reconfiguring CalMesh boxes, were deployed at the scene. The VPN server was located at UC San Diego. Five VPN nodes were built with cameras hardwired as clients. Port forwarding was then set up on the VPN nodes to enable connectivity (and therefore access to the cameras from other VPN clients). Several Windows and Linux laptops were used for this. Four of the VPN nodes were set up to use the Tropos network; the fifth was set up to use the CalMesh network. Internet connectivity was provided by a service provider, SkyRiver Communications. The Tropos gateway node was connected to the SkyRiver network via point-to-point microwave link. Only the Tropos network remained installed at the scene, all other parts were mobile and thus, uninstalled following the Mardi Gras event.

For a variety of reasons the current installation of the Tropos network in the GLQ is not optimal for ongoing measurements and testing during non-event periods (i.e., day-to-day testing). Therefore, the network is currently accessible only by password, which reduces the traffic substantially (and there is no 'ordinary' traffic of the type we wish to assess on the network). Issues include local businesses and hotels providing Wi-Fi to paying customers (competition concerns) and the City of San Diego is having significant electrical problems in the downtown area. The Gaslamp Quarter is an historic part of downtown with an aging civil infrastructure on which they are working. We are currently drawing up plans to develop an additional venue outdoors where such testing can take place, as well as working to resolve problems in the Gaslamp to enable a reliable testing platform situation. It is noteworthy to point out that the barriers which interfere with ongoing, extended testing of regular traffic flows would not exist in a real emergency or disaster.

The performance of Wi-Fi mesh networks will be affected by interference in dense urban areas (ISM band). Also it observed some portion of cellular spectrum will interference with Wi-Fi mesh nodes. Guarantee quality of service for streaming video from security cameras will dramatically affected after three hops. Considering all the effects, GLQ test bed deployment showed promising results. With a fixed mesh network on top of the lamp posts (Tropos node) a relatively reliable backhaul was provided for portable CalMesh nodes. CalMesh nodes are used for last hops extension of the network in the situations that fast deployment in a temporary base is a must.

Power for fixed mesh nodes is not available on top of all lamp posts in downtown areas. This caused lots of problems for deployment; most of the older version lampposts have a centralized light sensor. Therefore during the day there is no power on top of the posts.

CalMesh nodes are small, lightweight, and easily reconfigurable nodes that quickly self-organize to form a reliable wireless mesh network in any environment. CalMesh nodes use IEEE 802.11b-based Wi-Fi technology for users to easily set up communication during emergencies. The intelligent gateway nodes communicate with each other using an advanced wireless multi-hop relaying process. CalMesh nodes can use any of a variety of wired and wireless backhauls, which are integrated with the Always Best Connected system to allow users to seamlessly roam across and aggregate multiple network infrastructures and access technologies to maintain constant and optimal connection to the Internet.

UrbanMesh, a new version of CalMesh, with a more robust disclosure and integrated GIS capability will be ready within next 3 months. It will have up to 12 hours integrated battery. Within the next 6 months, the GLQ coverage will be expanded to double by using UrbanMesh nodes. UrbanMesh nodes are more compact and have sophisticated power management system; there is no need for coordination with city for power supply. Also, the added GPS capability will provide the ability of an advanced GUI control environment for monitoring and management of the nodes remotely.

UrbanMesh nodes will deploy in outdoor environments for a relatively long period of time. This can cause heating problem for boards, since ventilation is not possible. Different heat-sink mechanisms have been tested in the labs. But it can cause some problem in the real deployment. This issue needs more investigation and test, which will be conducted during next year.

By opening the network to public, GLQ test-bed can be used to test other community based application developed within RESCUE and other projects. This will enable different research teams to deploy and test their prototype under a fully loaded network, with realistic traffic patterns.

CalMesh: A Portable, Wireless, Ad Hoc, Mesh Network

CalMesh is an ad-hoc network of small, lightweight, and easily reconfigurable nodes that quickly self-organize to form a reliable wireless mesh network. CalMesh is designed to be rapidly deployed at the site of a crisis to restore the communication fabric crucial to emergency response, and has been used in conjunction with drills of the San Diego Metropolitan Medical Strike Team (MMST), a team of local, state, and federal responders who work together to develop and implement response plans for major urban crises and disasters.

CalMesh Technical Specifications

Wireless LAN Specifications:
- Standards-compatible: 802.11b, 2.4GHz, WDS
- Operating range: 300m outdoors
- Tx power: 200mw
- Antenna: 5dBi

Wide Area Access Specifications:
- 1xRTT, 1xEvD0
- GPRS/EDGE/UMTS

LAN Specifications:
- Auto-sensing 10/100BaseT Ethernet
- 802.11 bridging with STP

Management:
- Web based statistics capture
- Management over secure-shell interface
- Per node protocol-packet traffic analysis
- Zero configurations

Networking:
- DHCP/DNS/NAT
- Automatic mesh formation
- Configurable metric for the mesh nodes path-cost
- Internet Gateway node discovery mechanism
- Load balancing between multiple Internet gateways
- 802.11b access point capability

Triangulation:
- GPS support

Hardware:
- Weatherproof Aluminum Case with Motherboard featuring plug-ins for WLAN 802.11a/b/g, 1xEvD0 & 1xRTT, GPRS/EDGE, UMTS, WiMAX, and GPS
- Weatherproof Multiple Antenna ports supporting multi-band wireless communication, 800MHz, 900MHz, 1800MHz, 1900MHz, 2.4GHz, 3.5GHz, 5.2GHz, & 5.8GHz operation, with omni-directional or high gain antennas
- Weatherproof Ethernet and Serial ports for wired communication with PDAs, PCs, Handheld GPS, Bull Horns and other peripherals.
- Weatherproof Multimodal power system with 10Ah 12Vdc NiMH battery pack with built-in charger. Input and output power to/from Laptop PC, Solar Panel, Flashlight, Battery Powered Tools, Car Battery, or other available power sources.

The inside of a CalMesh node

CalMesh nodes use IEEE 802.11b-based Wi-Fi technology for users to easily set up communication during emergencies. The intelligent gateway nodes communicate with each other using an advanced wireless multi-hop relaying process. CalMesh nodes can use any of a variety of wired and wireless backhauls, which are integrated with the Always Best Connected system to allow users to seamlessly roam across and aggregate multiple network infrastructures and access technologies to maintain constant and optimal connection to the Internet.

Thirty-two mesh nodes were fabricated at UCSD and have been used in different drill activities throughout San Diego County. Several nodes were deployed with other ResponSphere partners, notably, the City of Champaign and UCI, for testing in various use-case scenarios and obtaining user feedback.

The City of Champaign recently took delivery of an Illinois Transportable Emergency Communications System (ITECS – communications trailer, see pictures below) that will be used within a 14 County area established by the Illinois Emergency Management Agency (IEMA) for disaster recovery purposes. Further, the units may be used statewide or interstate as directed by the state's Emergency Operation Center (EOC). If the City had the trailer last year, it likely would have been deployed to Louisiana or Mississippi.

Folded Position

Deployed Position

The ITECS unit houses communications equipment kept operational by onboard fuel and generator. The unit provides multiple forms of communication including various band radios and repeaters, laptop computers, telephones, PA system, satellite communications to the state (EOC), and satellite Internet access. The City will integrate the 3 CalMesh wireless network nodes into the ITECS unit. This will extend wireless communications to mobile devices beyond ½ mile (local) range of the trailer, either along or around a catastrophic event. Staff will then be able to use mobile computers for applications and Internet access for sharing information regarding the incident.

The first walk-through test of the ITECS unit will be early June followed by deployment tests every 30 to 45 days. Regular tests are required along with making a connection to the state's EOC otherwise the ITECS unit looses its satellite link.

The 3 CalMesh wireless network nodes will be put to good use and tested regularly. There are a number of linear and spatial patterns that the City's Information Technologies Department needs to test using the 3 CalMesh units assign to the City (in conjunction with the ITECS unit).

Additionally, the State of Illinois will eventually deploy 9 or 10 of the ITECS trailers, none of which are initially equipped with CalMesh units. I also understand there are other states considering similar mobile communication projects and devices, which means that the CalMesh artifact certainly has a market in emergency preparedness, but could be used for other less critical applications.

The researchers at UCI will use the CalMesh nodes for a number of crisis response scenarios both on an off campus. Disaster response events within the UCI Smart-Space will use the CalMesh routers to extend the Wi-Fi bubble to geographic areas that lack in 802.11 coverage. This can either be performed as a manual process by First Responders, or the Autonomous Sensing Platform (ASP) will drop these routers in a bread-crumb fashion. The various Wi-Fi technologies developed by the Responder Team (e.g., the EvacPack) will then be provided a data communications backhaul.

In scenarios where there is no Wi-Fi backhaul to extend, the CalMesh router can be equipped with an EVDO card to provide lower-bandwidth data communications backhaul and an 802.11 bubble. The Quality Aware Sensing Architecture (QUASAR) team will attempt to adjust the sensor output of any 802.11 device to match available backhaul bandwidth.

As an extension of the CalMesh project, Calit2 researchers are working to develop an easy and autonomous way to deploy the CalMesh network in outdoor environments. Gizmo Truck is our solution to this problem. The main goal for this project is to be able to send GPS final destination info to several Gizmo Trucks over the net, have them move to those locations autonomously and create the wireless network (CalMesh project). The number and/or the type of features which can be added to these Gizmo Trucks depend purely on the applications they will be used for. In fact, Gizmo Truck can be seen as a platform on wheels to deploy different technologies, not only a wireless network. The basic features installed in a Gizmo Truck are:

-Full motor capabilities, forward, reverse, and break.
-Web camera
-Software and web user interface.
-Override circuit for manual remote control.
    -Distance/Collision system

We built the first Gizmo truck with all the basic features mentioned above except the Distance/Collision system. This feature has been designed and it is fully functional but it hasn't been integrated in the system (we want to upgrade the truck chassis before). Also, the Distance/Collision software needs to be developed further to be able to respond to different obstacles and environment situations. The main problem that we are facing at the moment is noise interference which causes malfunctions in the manual override system. We tried to reduce the noise with some simple circuitry but it showed no improvement. Next step will be to upgrade the transmitter and receiver (which we believe are the main cause of the noise) and to install filters to get rid of the noise component in the signals.

In three months we plan to have 5 functional and chassis-upgraded Gizmo Trucks. Additionally, we plan to have the web user interface upgraded, revised and supporting more features such as new ways to remotely control the truck's motor and steering, and more reliable video stream capabilities. We will also develop Distance/Collision software adaptable to different environments.

We have plans to use these trucks as a platform which other projects can use to deploy their technologies. Some groups which have expressed interest in using the Gizmo Trucks include:

Computer Vision:
Inside of buildings it can be difficult to use GPS to determine precise position for a robotic agent. Taking advantage of simple, user-provided, visual tags we propose to use computer vision to determine position information for the Gizmo truck. These tags can also contain other tags, information or instruction in them so that the truck can more robustly navigate its space.

Networking & RF Modeling:
In both inside or outside scenarios, the deployment of a homogenous and stable wireless network is crucial in emergency situations. Providing Gizmo Truck with netstumbler capabilities, a set of five trucks (with the appropriate learning and adaptive algorithms) could sense where the best location to deploy the network node is. Also, trucks could create signal strength mapping data that can help us to learn how to distribute wireless networks in a more efficient and successful way.

We have further customized the CalMesh platform for another application, known as Mushroom networking. Using open access resources in an area, Mushroom can provide distributed aggregation, spatial diversity and redundancy (and thus reliability).

A Mushroom box (Mush-Box) with a camera/VPN was deployed in the UCSD Tech Ops Center during the Mardi Gras exercise. The connection remained stable throughout the Mardi Gras event. A second Mush-box was deployed in the hallway, outside the Tech Ops Center. The two boxes formed a self organized mesh network performing distributed gateway bandwidth aggregation. Leveraging statistical multiplexing, each of the Mush-Boxes aggregates all the available Internet access resources in a distributed manner through multiple gateways and provides this aggregated peak bandwidth to a single user.

In this particular demo, both Mushroom boxes were associated to the same access point in the hotel. Various tests were performed demonstrating the spatial diversity feature of the Mush-Boxes. First over a single box, the downlink Internet access speed was tested; then the second Mush-Box is switched on. Each experiment is run for two minutes. The average downlink speed obtained from each test over a single Mush-Box was 61.91 KB/sec. The average downlink speed over two boxes was 91.225 KB/sec. Despite the fact that both boxes were associated to the same access point in the hotel, the overall downlink speed was improved by 50% due to the spatial diversity. At Mardi Gras, they could have been used as back-ups in case of a Tropos backhaul failure, which was not necessary.

RF and GIS Modeling

We have been collecting measurements using several different RF measuring tools to determine the requirements for the GIS modeling equipment that we intend to purchase.

Adaptive Cellular Networking System:
The main products created out of this project is a cellular simulator which can interact with other simulators for inputs such as mobility and damages to infrastructure in order to provide analytical outputs in terms of the wireless coverage and capacity for the system as a whole. This simulator is also a part of the METASIM simulator.

A prototype was developed for all four simulation modules, and serves as a test of many key issues, such as timing, file transfers, and the ability to call the various components as external modules. Detailed Investigation of Opnet capabilities is scheduled to be completed by end of June 06. This includes different interfaces and databases that can be useful for other simulators. The investigation has already started and few of the format and interfaces have been identified. A test simulation for UMTS network has been conducted and the capabilities of Opnet have been tested.

The theoretical research group investigated a new multidisciplinary research solution for using emotional content from the human speech for providing Quality of Service (QoS) in a wireless mesh network. As part of this research, we teamed up to investigate this problem. We investigated the possibility of detection of distress (distress is defined in this context as the emotion content such as fear) in human speech and developed a theoretical model for classifying the speech into either a distressed call or a normal call. Once a call is identified as a distressed call, the voice packets generated by that call are marked in order to indicate the source’s distressed status. The voice packets originated by a distressed or highly agitated caller are treated in a differentiated manner by the distributed wireless mesh network in order to provide better QoS. We developed a non-binary-based adaptive back-off mechanism to provision differentiated QoS for the voice packets generated by the distressed or highly agitated sources in comparison with the normal voice sources. The developed solution provides an average distress detection probability of 60% to 80% and an average end-to-end delay performance advantage of about 60% at very high load. With three traffic classes such as normal traffic, moderately distressed traffic, and extremely distressed traffic, the obtained accuracy of recognizing the emotion content in speech is about 60-75%. The emotion detection accuracy went up to 80-85% when there were only two classes for detection: (i) normal traffic and (ii) distressed traffic. Additional findings made by our experiments include that the use of a non-binary based adaptive back-off algorithm could provide about 60% reduction in the end-to-end delay for the calls originated by a distressed or highly agitated source. As part of the knowledge base creation, we conducted a detailed study on Issues and solutions in Wireless mesh networks and Load balancing techniques in wireless mesh networks. The findings from these studies are developed into two separate invited text book chapters.

In the next three months, our primary objective is to investigate the capabilities of the OPNET simulation platform for research into adaptive cellular networking system. The deliverables for the next 12 months include the following:

1. Integration of OPNET with other simulators.
2. Interaction with an external database, so that the simulation results can be read iteratively by other simulation functions.
3. Accept agent location from DrillSIM iteratively.
4. Specification of all input and output file formats as well as timing for information exchange.
5. Study of different crises situations and their influence on the cellular infrastructure.
6. Obtaining the location of cell sites (collaboration with carriers).
7. Integration of real-time measurements (TEMS) into OPNET-based cellular simulator

In the GLQ testbed, measurements of network traffic flows were taken (the VPN boxes were assigned static IP addresses in the Tropos system). Three of the nodes and the SDPD command post served as the main locations from which data samples were collected during the Mardi Gras celebrations. (Additional preliminary measurements were taken in the area during preparation for and installation of the Tropos network.)

Taking Measurements in the GLQ testbed

Measurement samples were taken intermittently from different locations in the Mardi Gras area during the festivities. Packet delays for the camera nodes to the VPN server were around 50-100ms. Throughput of the VPN connections going through the CalMesh network was around 200kbits/sec. Because the VPN runs on TCP, TCP traffic constituted a large amount of the entire traffic measured, 86.54%; UDP traffic accounted for 10.81% of the traffic. Sample results in bytes transferred at the 5th and E streets corner at 8:25 pm for 15 minutes were as follows: VPN node/camera IP address (number of bytes): 10.100.10.120 (10025033), 10.100.10.116 (895549), 10.100.10.118 (34808) and 10.100.10.112 (20428).

The highest observed traffic flow during the Mardi Gras originated at VPN box 10.100.10.120 (located at 5th and E streets) during a sampling period of about 900 seconds at approximately 10 PM. During the observation period, this VPN client had two separate connections in time. This may be due to a dropped first connection, followed by the initiation of another connection. The traffic variation with respect to time at 5th and G streets was rather different compared to that observed at 5th and E. Here, the traffic sample (5th and G) was stable without breaks. Also of note, for every TCP connection pair, a large bi-directional traffic difference was noted.

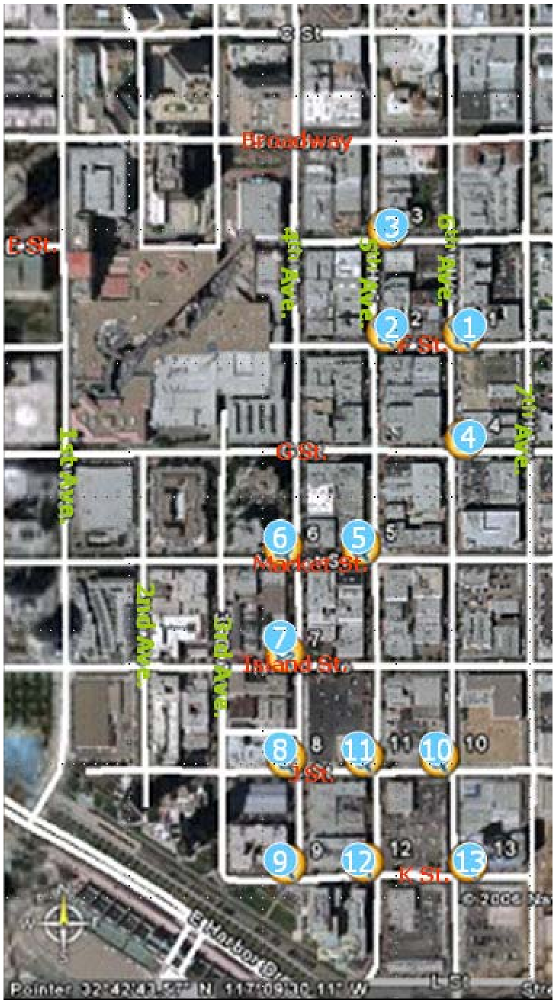

Map of NetStumbler Measurement Locations.

1. *Martini Ranch, 6th and F St.*
2. *The Bitter End, 5th and F St.*
3. *Ostera, 5th and E St.*
4. *Buca De Peppo, 6th and G St.*
5. *Whiskey Girl, 5th and Market St.*
6. *Starbucks Coffee, 4th and Market St.*
7. *Aubergine, 4th and Island St.*
8. *Jolt'n Joes, 4th and J St.*
9. *Fleming's Prime Steak House, 4th and K St.*
10. *Five Star Parking, 5th and J St.*
11. *Splash, 5th and J St.*
12. *Dussini Mediterranean Bisto, 5th and K St.*
13. *Marriot Parking Garage, 6th and K St.*

The UCSD Calit2 Responsphere team is creating a Wireless Communications Mobile Command and Control Vehicle. A detailed survey of existing command and control vehicles was performed. This activity included on site inspection of the existing vehicles and interviews with the personnel who use and deploy these vehicles. Most of these vehicles are provided by state and county police and firefighter units, with the remainder provided by Army Reserve and National Guard units. Most of the vehicles are underutilized, with an average of 1 or less real deployments per year. In fact, most of vehicles participate in more drills of about 3 to 4 per year compared with unexpected incidents. The exceptions are vehicles with satellite backhauls and Plain Old Telephone Service (POTS). These vehicles or trailers provide 10 voice lines and 2 56kBaud data lines, and are used more than 6-8 times per year in real events, of which about 3-4 are unexpected events. Many of the other vehicles suffer from severe data rate bottlenecks. The vehicles may be equipped with Ethernet routers with 1Gb/s total capacity, but local on-site ground communication is limited a few 2400-9600baud wireless modems, and a 128kb/s – 512kb/s satellite backhaul.

Some vehicles provide radio interoperability between police, firefighters, and paramedics, etc. The radio interoperability vehicles suffer from a poor interface establishing cooperative groups. For example, operator does not have any way to remotely download a policeman’s radio with a temporary list of firefighter radio addresses. The policeman has to listen to the operator’s voice instructions, and then push a lot of buttons on his radio to establish cooperative groups. In addition, with interoperability mode activated, all radio traffic has to be routed through the radio interoperability vehicle whether or not the police or firefighters need to talk to each other. This limitation can cause radio coverage problems. Most of these radios connect to a mountaintop base station or fire truck, and usually relayed to local headquarters. The interoperability radio vehicle may not be in a location that has coverage for the mountaintop base stations, fire trucks, or ambulances, etc.

Key information to a from the command and control vehicle is often disseminated to the local responders from a large white board hung on the side of the vehicles with maps laid out on tables under a tent. People come by with clipboards, talk to each, write important items down on their clip board, and return to action.

During our team’s visits to San Diego Engine Co. No. 35 and UCSD’s campus police department, that fire trucks and police cars lack radio relays. For example, the handheld radio must communicate directly with the remote mountain top base station, or directly to another handset. For example, a radio relay cannot be set up from a radio handset to the fire truck and then to the remote base station. Likewise, a radio handset cannot relay from to another handset and then to the remote base station. This capability is required when the office can make radio contact with another officer or his police cruiser, but not the remote base station. This spotty radio coverage occurs in underground parking garages, and in the shadow of numerous cliffs in San Diego. An example workaround is for a middle man firefighter to hold 2 radio handsets, and repeat what is spoken in both directions.

The limitations described above needs to be overcome. Consequently, we have developed new approaches that will lead to success at this time. Our wireless infrastructure relies on getting both the communication into an IP network as soon as possible. This includes both human users and wireless sensors. Consequently, even traditional Push-to-Talk voice communication will be adapted to VoIP. The use of mesh networks allows for seamless relays in the event of spotty wireless coverage. The firefighter for police officer does even need to be aware that relaying is taking place. The network layer autonomously decides to relay based on Always Best Connected (ABC) protocols developed at UCSD. Using VoIP also provides the user with Call to Collaborate features developed at UCSD that can create teams without each user having to program their radio in great detail.

A foldable flat panel display for outdoor viewing has been developed. The display consists of 4x3 array of 19” extra bright LCD screens that can be disassembled for portable storage. This display can augment the traditional whiteboard, but can be quickly be reconfigured based on the particular team meeting at the command and control vehicle. The traditional clipboards can be augmented by WiFi connected tablet PCs. Information can be transmitted and received by the Tablet PC as users enter information or pick up information from nearby wireless sensors. In this way every member of the team can act as information courier, even when moving in an out of wireless coverage. The command and control vehicle can be provisioned with about a dozen table PCs.

We will be completing construction of 5 more antenna caddies in the next 3 months. We are also developing a micro-gasoline powered generator for powering mesh network nodes, antenna caddies, and the satellite link for several weeks. Our new motor generator will be coupled with a rapid charge circuit to recharge Pb-Acid and NiMH battery packs in 30minutes, or Li-Poly battery packs in 1 hour. The motor generator will automatically turn on and off to maintain the batteries. The micro-gasoline power motor generator will be prototyped over the next 6 months.

The Wireless Communications Command and Control Vehicles will consist of a 35ft to 40ft long “Toy Hauler” 5th wheel trailer pulled by a crew-cab pick-up truck. This vehicle with benefit project RESCUE, WIISARD, and DARPA DSCR Wireless Infrastructure Reconstruction Project.

Besides communication equipment, the vehicle will contain a collapsible outdoor multi-monitor LCD display array. This display will be used to display large scale tables, figures, maps, and graphs that can be changed to suit the needs of the team. In addition, we will equip the command and control vehicle with a poster-size printer fed with special waterproof paper to create large maps and tables that team members can layout on tables, scribble on, and take out into the field. A smaller clear plastic laminator machine will also be available to rapidly create reference data cards.

The main technical challenge is developing a high data rate backhaul from the command and control vehicle to the internet. Satellite backhaul is the obvious choice, but the high cost and limited availability of > 2.048Mb/s data links is a major bottleneck. Cellular data services such at HSDPA and WiMAX promise data rates of greater than 10Mb/s, but these signals need to be accessed when out of coverage or nearby cellular basestations are down. Consequently, our team is developing methods to access distant cellular towers. Cellular antennas are often down tilted for coverage area control. However these signals can be accessed from hilltop receivers strategically placed to pick up the ground bounce, and retransmitted to the command and control vehicle.

HAM radio operators have developed low cost uWave and mmWave reflectors on mountaintops that can be used for opportunistic communications. In addition, many police departments, fire departments, and TV stations have their own passive reflector systems that could be used on a short term permission basis.

Long distance WLAN communication is possible from a UAV over the command and control vehicle to a distance receiver. The 802.11n standard has the required Doppler shift immunity and could be used to establish >10 Mb/s communications over 5 -15km.

Our will incorporation our mesh network bit-torrent like technology to increase bandwidth for data communications. We will aggregate several 128kb/s to 1.024Mb/s channels to create a mushroom network. This method works well for delay tolerant networks.

Police, Firefighters, and Paramedics could utilize this vehicle in both their drills, and long term deployments. For example, this vehicle could be utilized near forest fires lasting several days, or hurricane aftermaths lasting several weeks. With the increase in Homeland Security cooperative drills, we can place our Wireless Command and Control Vehicle in the hands of end users without commercial “buy it now” pressures. For example, instead of UAVs, we could utilize the Civil Air Patrol’s light aircraft to act as a communications relay for Red Cross logistics support. Our test bed will encourage criticism and feedback from multiple agencies without intimidation. Being a university we are a neutral party that is not a threat to any agencies’ jurisdiction.

The main educational outcome is the training of ECE, MAE, and CSE students in sub-system and system integration. ECE, MAE, and CSE students will work as team with cooperative mission that will allow the students to discover their strengths and weaknesses.

The main deliverable is the Wireless Communications Command and Control Vehicle. The vehicle provides concise objective within a limited scope. The vehicle functions as a living test bed that is expect to have a 5 to 10 year life span.

The intended audience includes but is not limited to government agencies in charge of coordinated emergency response management. Primary examples include the Department of Homeland Security (e.g. FEMA), California Emergency Management Division, and the San Diego Emergency Response Council.

Technologies developed to date for the vehicle include (1) A WiFi Public Address System for Disaster Management, (2) Mesh Network Antenna Caddy, and (C) Flexible Satellite Deployment System, all of which are described in more detail below:

(1) A WiFi Public Address System for Disaster Management

In disaster settings sometimes shouting is the best form of communications. If background noise levels are impractically high, amplification often via “bull horns” is proven solution. However, sometimes at disaster sites, one might need to make announcements or broadcast instructions over a very large area without any kind of pre-existing public address system. As part of the Wireless Internet Information System for Medical Response to Disasters (WIISARD) we developed a WiFi disaster bullhorn broadcasting system. The unit used the mesh 802.11b network deployed at the scene to create ad hoc wide area public address system.

The WiFi Bullhorn is wireless Linux computer coupled to a sound amplifier, a batter source and two "bullhorn style speakers. It is controlled by a web page served by the device. The web page offers stored messages (Eg. "Please proceed to the nearest exit."), warning alarms, soothing music, as well as the ability to record and broadcast custom messages.

Deployed WiFi "Bullhorn" unit

The web server is Apache2 with PHP4 and the web server controls am audio player (mpg123); the operating system is Debian GNU/Linux running a custom 2.6 kernel. The kernel runs on a Soekris net4521; it is essentially a motherboard with a 133MHz 486, 64MB RAM, two PCMCIA slots, a CompactFlash slot, and a type III PCI slot. The operating system was stored on a 512MB CompactFlash card which was set to read only to allow quick power downs without damaging the file system. This posed as a challenge for storing custom messages; we solved this by writing custom messages to a tmpfs (RAM disk) at the cost of being limited to one custom message at a time which is lost when power is cut or a new message is recorded. The soundcard used is the MP878-S by Commell; We used off the shelf 802.11b PCMCIA cards to connect to the network.

Since rapid deployment is essential in disasters, a design requirement for the WiFi bullhorn was that it would be light and simple to set up. We used a standard aluminum engineering tripod as the base. The bullhorn chosen is the BDT30A by Bogen Horns. The bullhorn uses a Jenson UPA224CS amplifier which is lightweight and could fit in the same container as the Soekris board and battery (a 10,000mAh NiMH 12V battery). In addition to being lightweight, the unit needed to be weatherproof, so damageable parts were put into a waterproof Pelican case. The case was retrofitted with weatherproof audio jacks and antennas, as well as a power switch and LED's.

The device was tested at a Southern California Disaster Drill held at the Del Mar Fair Grounds. This event simulated explosion of a large car bomb with resulting a pesticide release. Contaminated victims were rescued by emergency response-teams. Two bullhorn units were deployed; one at the location of the bomb (to potentially play a warning message about contamination), and the other near the treatment areas (about a kilometer apart). At the drill we received positive feedback from emergency workers who generally said they could see the potential uses for such a device. There were some limitations resulting from a design that required pre-recording and transmission, rather than streaming of audio messages.

(2) Mesh Network Antenna Caddy

The focus of this project is to create a way to use CalMesh network boxes in a wider range of applications. In order for the mesh network to function, the antennas must have a good line of sight to send and receive signals. This was achieved in the past by placing boxes on the top of buildings, or other locations that happen to get good signals. However, this is not always possible, and the existing mesh network boxes were therefore limited in their functionality. Also, the boxes used battery packs which needed to be charged after less than 12 hours of use. We built a system that could be set up virtually anywhere, and that could sustain itself for a matter of days.

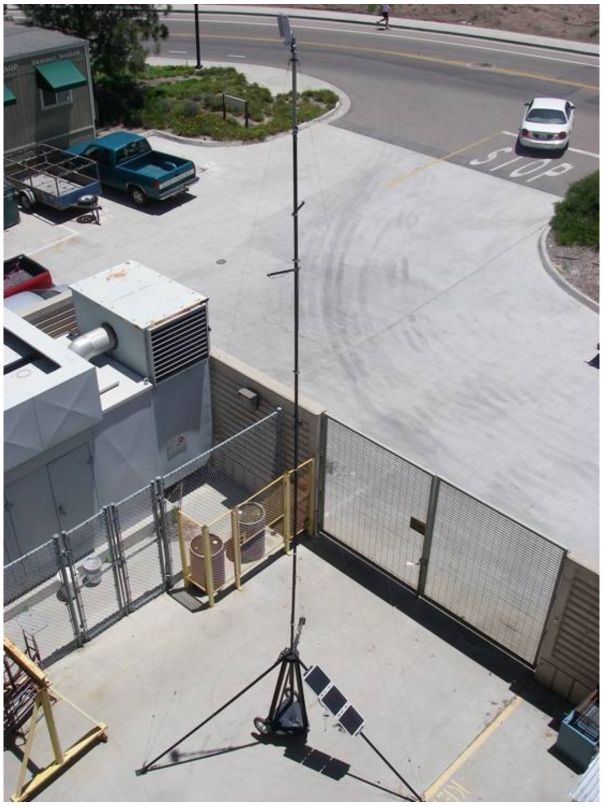

Deployed Antenna Caddy

The mesh network antenna caddy is comprised of a segmented mast which is supported by a tripod, all centered on a rolling base large enough to hold all the necessary equipment and any future additions or modifications.

The caddy may be transported in two main sections, each weighing under 23 kg (50 lbs). Once near the job site, the caddy is partially assembled into a rolling structure that may be moved easily by one person. Air filled tires allow the caddy to be pulled over rough terrain, and all sensitive equipment and electronics are kept in a watertight case to allow submersion if necessary.

Once at the site, the caddy may be assembled by one person to a height of up to 9.14m (30 ft). The mast is supported by two sets of guy wires that attach to the ends of the three legs of the tripod. Guy wires are also used to control the spread of the legs from the base. The result is a sturdy, lightweight structure that does not need to be staked to the ground to be stable. It does also have the option of stakes through the feet, if the conditions are especially hazardous.

The caddy may be configured in multiple ways, and has adjustments at virtually every level. The segmented mast allows for multiple mast heights, depending on the line of sight needed. The directional antenna is mounted on an adjustable mount, which allows for a down tilt, and the mast may be rotated from the base to make fine adjustments. The antenna mounts and guy wire mounts clamp to the mast, and may therefore be placed anywhere along the height. The length of the tripod legs is also adjustable, which allows for uneven terrain or a variable footprint radius.

The circuitry is powered by two 12V, 18aH deep cycle batteries, and by three Coleman solar panels. The batteries and mesh network board are kept in a watertight enclosure, with bulkhead connectors which may be replaced if damaged. The solar panels and batteries are connected so that the panels charge the batteries when possible, but the network will still function at night or low light. Even without solar power, the batteries are adequate to provide power for up to eight days of normal network use, with solar power extending this deadline indefinitely.

(C) Flexible Satellite Deployment System

We have also constructed a flexible deployment system for our ViaSat LinkStar satellite data link. The current mounting scheme allows for deployment from or on top of a truck, or on the ground. In this way, the satellite link can be decoupled from the vehicle, and left in a better location with a view of the southern sky. In some cases, the vehicle may need to park in a location without a southern sky view. We have found this to be true in urban canyon environments. Photographs of the flexible deployment satellite link are show below. The dish diameter of 1.8m is sized for antenna gain so that only a 2W RF amplifier is required for a 512kb/s uplink. This low power RF amplifier with a high gain parabolic dish increase usable time in battery powered applications.

Satellite deployment system (stand-alone)

Satellite deployment system as used in vehicle at Mardi Grad exercise, 02.28.06

Portable Tile Display Wall for Visualization in Crisis Response

To date we have begun to assess which technologies will be most easily integrated into the Tile Display Wall, and which will meet the stringent requirements (particularly, robustness of hardware and brightness of display technology) for crisis response applications. We have begun sampling different pieces of hardware (in particular, tablet PC laptops) to see which will be most easily integrated into the complex hardware-software visualization solution we are developing. One of our biggest challenges will be to develop a visualization system that is simple to use yet meets the complex display requirements of the end users. We will address this through trials. We plan to have an initial trial of a prototype prepared by the MMST drill training in July, which will leave us time to work out some of the kinks in order to complete the second phase of the project by the drill on August 22nd. Though no specific products have been developed to date, we are currently developing this display wall, which will include both the hardware portion (a uniquely mounted system of disparate display technologies which meets the portability and robustness requirements of the application) and the software portion (a complex exercise in systems integration built on a Linux platform).

Anticipated deliverables for the next 3 months are briefly described above. After the MMST drill in August, we plan to continue to improve the software system for more seamless integration of technologies over the course of the next year, as we will have developed a better initial understanding. This is a complex exercise to create a cohesive system out of components that were each intended for stand-alone use. Therefore, getting them to communicate cross-platform with different applications is no simple task. However, we are hopeful that with regular interaction and frequent iteration with our possible end-user community (members of the RESCUE and WIISARD projects, as well as the response organizations these projects work with including as fire, police, medical, etc), we will be able to develop a straightforward and cohesive system. There are no specific educational outcomes and deliverables planned, however, there is the potential to integrate this into ongoing ResponSphere educational activities, and it will become a platform for future development which can be used as a basis for ECE 191 and other project courses. No system like this has been developed in the past; therefore we plan to present the results of our research at various relevant conferences organized by IEEE, and other such organizations, as well as venues such as the Hazards Center Workshop, and any conferences or workshops that use UCSD Calit2’s Atkinson Hall as a host-facility. Future team members will depend on the applications to be developed on this platform, and will likely include a subset of current RESCUE, WIISARD, and ResponSphere project members, partners, and collaborators.

Currently, the cost of high-resolution visualization displays is prohibitive to universal adoption of this technology. The only permanent installation of a 4K projector in the nation is in the Calit2 Digital Theatre in Atkinson Hall at UCSD. Though breathtaking, the 4K resolution (only about four times the resolution of HDTV) is still not sufficient for many applications which require visibility of extremely detailed imagery. To adapt such high-resolution technologies for outdoor use is even more cost-prohibitive, as cost increases drastically for large displays with sufficient luminosity to be viewed outdoors. By combining multiple tablet PCs with outdoor-viewable screens, the system is constructed at a fraction of the cost. These tile displays are becoming an increasingly popular way to visualize large amounts of high-definition data; however our approach is unique for several notable reasons. First, the high-luminosity for outdoor applications; second, the display is being made up of screens of varying sizes (12” up to 32”) and resolutions; third, each of the tablets is being configured to function both as a member of the tile display and independently, so that they can be pulled off the panel to run individual applications.

In addition to the display wall itself, we have developed several applications and services which are eager to utilize this technology. Those include (A) MoVs: Mobile Vision System, (B) LBT: Location Based Tracking, (C) Peer to Peer Information Collection System, and (D) ESB for System Integration.

(A) MoVs: Mobile Vision System

MoVs is a general purpose, low-cost, mobile system geared for computer vision applications which we are developing. The MoVs platform is built on top of the TI OMAP5912OSK development kit, running open source linux. The goal is to implement the platform in such a way that porting algorithms is as painless as possible, supporting the workflow of Matlab development & Debugging of original algorithm, as well as development in OpenCV (or other open source library) of application on desktop PC platform, and recompilation inside of MoVs environment and embedded debugging.

There are several thrusts of this work: Development of build environments and techniques to make the porting process painless, camera driver development and implementation, and Implementation/Optimization of the open source libraries, either through the on-board DSP or through external FPGA development.

Some of the key components include:

1. Mobile (embedded) implementation of the License Plate Recognition System
2. MoVs board – Texas Instruments OMAP 5912 devlopment kit outfitted with Linux, drivers, and the OpenCV libraries to support embedded implementations of Computer Vision software.
3. MoZi prototype: collaboration with the ZigZag project has produced a two-servo system to be driven by the mobile vision project for guiding the visually impaired.
4. MobileVision / Mesh prototype: Mobile Vision board embedded in a calmesh mesh network node, to add embedded computer vision capability to camera- enabled meshnodes.

Moving forward, the primary adjustment to the strategic timeline is that temporarily, application development is suspended until the platform can be finalized and delivered. Then the focus will shift away from development of the platform to development of the applications and optimization of the platform. In 3 months the platform will be completed, along with documentation (online) for developing for the platform. The first application-projects should have been started. In 6 months, the first prototypes of the applications should be implemented. In one year, the first round of applications should be nearing completion, and the hardware platform review should have been conducted to motivate redesign/evolution of the project (which will be the next phase after that point). We have a major presentation at the National Federation for the Blind conference in July. Potential end-users include the surveillance-related community and the visually-impaired/blind community. We are also developing a technical project narrative which will provide a design history of the project, as well as detailed instructions on pitfalls and practice of developing for and on our platform.

(B) Location Based Tracking System (LBT)

The mobile based tracking system is built with the latest Assisted GPS (AGPS) technology.Based on AGPS, a mobile phone is used to track the position and speed of objects. A real-time map is provided to view location, speed, and identity of objects. The system consists of a mobile client and a tracking/mapping server. Both client and server can work together as complete solution for object tracking. The GPS data used by our system includes the latitude, longitude, horizontal speed, heading, and altitude. The system will be integrated with Microsoft Map Point mapping engine to display a more sophisticated street map. The system can also work with a third party's mapping server as long as the server complies with the XML interface. The system is on 24x7; therefore it can be used anytime.

(C) Peer to Peer Information Collection System

We have developed a fully automated peer to peer system that can collect and relay disaster related information to the general public and to the first responders. Though government agencies and the private sector have some of the basic data needed for effective disaster prevention and management, the means to effectively disseminate the data in an intelligent manner (i.e, delivery of relevant and timely information to the right target population) is lacking. Typically the data is disseminated in a broadcast mode, which could create a mass panic type of situation. Also, in many situations, there is significant lag in the collection of crisis related data by the government agencies. This lag can be eliminated by empowering the general public to report relevant information.

We have used San Diego as a test bed to develop, deploy and test the above mentioned system to empower the general public (in particular the commuters) of the county 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 will be based on speech recognition. We have learnt 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). The system has should be available on a regular basis, disseminating information that is valuable to public on an every day basis.

We have addressed these problems by using a traffic notification system that has been operational for the past two years and used by thousands of San Diego commuters every day as the basis for prototype. The system currently provides personalized real time traffic information to the commuters via cell phones (http://traffic.calit2.net). 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. Also based on the severity of the incident, we can notify all or part of the users via voice calls and text messages in a parallel and scaleable manner. We will create a hierarchical voice user interface that will accommodate for the severity of the incidents being reported. Examples of scenarios are the following. In the simplest case, a commuter might see a major accident that has closed several lanes of a highway. He can report this incident via the system and other users who are calling in for traffic information will hear this event if it happens 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 all the users, to avoid that region of the freeway.

The traffic system is being expanded to cover Los Angeles and Orange counties freeways that have inductive loop sensors embedded in them.

(D) ESB for System Integration

In the system integration project, several important factors affecting project integration were learned. The RESCUE team utilizes a standards-based approach using an Enterprise Service Bus (ESB). The purpose of the ESB is to facilitate application and process integration by providing distributed processing, intelligent routing, security and dynamic data transformation. In an ESB these services are infrastructure services so each application does not have to implement these requirements independently and in a proprietary manner. The ESB addresses the disadvantages of existing solutions by creating a standard infrastructure for integration. Point-to-point solutions, where each of n components requires n-1 interfaces for full communication, are replaced by a bus solution where each component requires a single interface to the bus for global communication. An ESB provides distributed messaging, routing, business process orchestration, reliability and security. It also provides pluggable services and, because of the use of standards, these pluggable services can be provided by third parties and still interoperate with the bus.

The main research directions of the system integration project are (i) an Enterprise Service Bus (ESB) for integrating the different networking modules in RESCUE and (ii) a sophisticated web portal-based front-end for the Enterprise Service Bus. This website can be accessed at www.dsscc.com. This site was developed to serve as a front end for the Enterprise Service Bus.

Products.

1. R.B. Dilmaghani, B.S. Manoj, B. Jafarian, R.R. Rao, “Performance Evaluation of RescueMesh: A metro-Scale Hybrid Wireless Network,” Proceedings of WiMesh Workshop, IEEE Workshop on Wireless Mesh Networks, held in conjunction with SECON, Sep. 2005.

2. R. B. Dilmaghani, B.S. Manoj, R.R. Rao, “Emergency Communication Challenges and Privacy,” Proceedings of the 3rd International ISCRAM Conference, May 2006.

3. R. B. Dilmaghani, R. R. Rao, “On Designing Communication Networks for Emergency Situations” Proceedings of International Symposium on Technology and Society (ISTAS), IEEE, June 2006.Larry Smarr, Maxine Brown, Tom DeFanti and Cees de Laat (guest editors), Special issue on the International Grid (iGrid) 2005 Workshop, Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., to be published (Summer 2006)

4. Tom Peterka, Daniel J. Sandin, Jinghua Ge, Javier Girado, Robert Kooima, Jason Leigh, Andrew Johnson, Marcus Thiebaux, Thomas A. DeFanti, “Personal Varrier: Autostereoscopic Virtual Reality Display for Distributed Scientific Visualization,” Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., to be published (Summer 2006)

5. Jason Leigh, Luc Renambot, Andrew Johnson, Donna Cox, Jonas Talandis, Tom Coffin, Brian Davis, Paul Wielinga, Gee Bum Koo, Jae Yoon Kim, Brian Corrie, Pierre Boulanger, Byungil Jeong, Ratko Jagodic, Nicholas Schwarz, Dmitry Svistula, Rajvikram Singh, Julieta Aguilera, Xi Wang, Venkatram Vishwanath, Brenda Lopez, Daniel J Sandin, Tom Peterka, Javier Girado, Robert Kooima, Jinghua Ge, Lance Long, Alan Verlo, Thomas A DeFanti, Maxine Brown, Robert Patterson, Joe Reitzer, Tom Prudhomme, Bram Stolk, Todd Zimmerman, Sang Woo Han, Jong Won Kim, “The Global Lambda Visualization Facility: An International Ultra-High-Definition Wide-Area Visualization Collaboratory,” Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., to be published (Summer 2006)

6. Takashi Shimizu, Daisuke Shirai, Hirokazu Takahashi, Takahiro Murooka, Kazuaki Obana, Yoshihide Tonomura, Takeru Inoue, Takahiro Yamaguchi, Tetsuro Fujii, Naohisa Ohta, Sadayasu Ono, Tomonori Aoyama, Laurin Herr, Natalie van Osdol, Xi Wang, Maxine Brown, Tom DeFanti, Rollin Feld, Jacob Balser, Steve Morris, Trevor Henthorn. Greg Dawe, Peter Otto, Larry Smarr, “International Real-Time Streaming of 4K Digital Cinema,” Future Generation Computer Systems/The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier B.V., to be published (Summer 2006)

7. Yukio Tsukishima, Akira Hirano, Naohide Nagatsu, Takuya Ohara, Wataru Imajuku, Masahiko Jinno, Yoshihiro Takigawa, Kazuo Hagimoto, Luc Renambot, Byungil Jeong, Jason Leigh, Tom DeFanti, Alan Verlo, Linda Winkler, “The First Application-Driven Lambda-on-Demand Field Trial over a US Nationwide Network,” Proceedings of OFC/NFOEC 2006 (Optical Fiber Communication/ National Fiber Optic Engineers Conference), sponsored by IEEE/ComSoc, IEEE/LEOS and the Optical Society of America (OSA), March 5-10, 2006, Anaheim, CA, paper PDP48,

8. Daniel J. Sandin, Todd Margolis, Jinghua Ge, Javier Girado, Tom Peterka, Thomas A. DeFanti, “The Varrier Autostereoscopic Virtual Reality Display,” ACM Transactions on Graphics: Proceedings of ACM SIGGRAPH 2005, Los Angeles, CA, July 31- August 4, 2005, Vol. 24, No. 3, pp. 894-903

9. J.P. Schulze, A. Chourasia. "A User Interface for High Dynamic Range Transfer Function Design." Poster to appear in DVD proceedings of ACM SIGGRAPH 2006, Boston, MA

10. Akin, O., Ergüt, S., and Rao, R. “Client Side Active Queue Management for 3G Cellular Networks,” Proceedings of IEEE Consumer Communications and Networking Conference, 01/2006.

11. Arisoylu, M., Mishra, R., Rao, R., and Lenert, L. “802.11 Wireless Infrastructure To Enhance Medical Response to Disasters,” Procceedings of American Medical Informatics Association 2005 Annual Symposium, Washington, D.C., 10/22/2005.

12. Braunstein, B., Mishra, R., Manoj, B.S., and Rao, R. “On The Performance of Distributed Wireless Mesh Networks for Disaster Response,” IEEE MDC, 2006.

13. Braunstein, B., Trimble, T., Mishra, R., Manoj, B. S., Lenert, L., and Rao, R. “Challenges in Using of Distributed Wireless Mesh Networks in Emergency Response, Proceedings of the 3rd International ISCRAM Conference, 05/2006.

14. Chebrolu, K., Raman, B., and Rao, R. “A Network Layer Approach to Enable TCP over Multiple Interfaces,” ACM/Kluwer Journal of Wireless networks (WINET), Vol. 11, No. 5, 09/2005.

15. Chebrolu, K. and Rao, R. “Bandwidth Aggregation for Real – Time Applications in Heterogeneous Wireless Networks,” IEEE Transactions on Wireless Communications, 2006.

16. Dilmaghani, R. B., Manoj, B. S., and Rao, R. R. “Emergency Communication Challenges and Privacy,” Proceedings of the 3rd International ISCRAM Conference, 05/2006.

17. Dilmaghani, R. B. and Rao, R. R. “On Designing Communication Networks for Emergency Situations,” International Symposium on Technology and Society (ISTAS '06), June 9-10, 2006.

18. Dilmaghani, R. B., Manoj B. S., Jafarian, B., and Rao, R. “Performance Evaluation of RescueMesh: A Metro-scale Hybrid Wireless Network,” Proceedings of WiMesh Workshop, IEEE Workshop on Wireless Mesh Networks, held in conjunction with SECON, 09/2005.

19. Hegde, R., Manoj, B.S., Rao, B. and Rao, R. “Emotion Detection from Speech Signals and its Applications in Supporting Enhanced QoS in Emergency Response,” Proceedings of the 3rd International ISCRAM Conference, 05/2006.

20. Hegde, R. and Rao, B. “Blind Single Channel Speech Segregation using Sinusoidal plus Residual Modeling,” Technical Report, 2006.

21. Lenert, L. A., Palmer, D. A., Chan, T. C., and Rao, R. “An Intelligent 802.11 Triage Tag For Medical Response to Disasters,” Procceedings of American Medical Informatics Association 2005 Annual Symposium, Washington, D.C., 10/22/2005.

22. Manoj, B. S., Mishra, R., and Rao, R. “SEBAG: A New Dynamic End-to-end Connection Management Scheme for Multihomed Mobile Hosts,” Lecture Notes in Computer Science (LNCS), Vol. 3741, 2005

23. Manoj, B. S. “On The Use of Bandwidth Aggregation Over Heterogeneous Last Miles,” Proceedings of International Conference on Broadband Networks, 2005.

24. Rossi, M., Rao, R., and Zorzi, M. “Cost Efficient Routing Strategies over Virtual Topologies for Wireless Sensor Networks,” Proceedings of 2006 IEEE Globecom Conference, St. Louis, MO, 11/28/2005.

25. Rossi, M., Zorzi, M., and Rao, R. “Statistically assisted routing algorithms (SARA) for hop count based forwarding in wireless sensor networks,” Wireless Networks, 06/2006.

26. Manoj, B. S. and Rao, R. “Issues and Challenges in Wireless Mesh Networks,” Book Chapter in Wireless Mesh Networks: Architectures, Protocols, and Standards, Edited by Y. Zhang et. al., CRC Press, 2006.

27. Manoj, B. S. and Rao, R. “Load Balancing in Wireless Mesh Networks,” Book Chapter in Wireless Mesh Networks: Architectures, Protocols, and Standards, Edited by Y. Zhang et. al., CRC Press, 2006.

28. Daniel Massaguer; Chien-Liang Fok; Nalini Venkatasubramanian; Gruia-Catalin Roman; Chenyang Lu. "Exploring Sensor Networks using Mobile Agents" In Proceedings of AAMAS’06 May 8–12 2006, Hakodate, Hokkaido, Japan, 2006-05

29. Xing, B., Deshpande, M., Venkatasubramanian , N., Mehrotra, S. "Revisiting Network Wide Broadcast in Ad Hoc Networks for Application Data" Technical Report, 2006-04.

30. Xingbo Yu; A. De; Sharad Mehrotra; Nalini Venkatasubramanian; N. Dutt. "Energy-Efficient Adaptive Aggregate Monitoring in Wireless Sensor Networks". Technical Report, 2005.

31. Xingbo Yu; Sharad Mehrotra; Nalini Venkatasubramanian. "SURCH: Distributed Query Processing over Wireless Sensor Networks"Technical Report, 2005.

Contributions.

1. CalMesh (more info at http://mesh.calit2.net/rescue/wiki/Main_Page and http://calmesh.calit2.net)
2. Portable Tile Display Wall for Visualization in Crisis Response
3. Mesh Network Antenna Caddy (more info at http://maeprojects.ucsd.edu/mae156/ws2006/A1/index.htm)
4. A WiFi Public Address System for Disaster Management
5. Flexible Satellite Deployment System
6. MoVs (more info at http://rescue.calit2.net/mobilevision , http://vision.ucsd.edu/lpr , and http://vision.ucsd.edu/collaborate/index.php/MoVs_Development_Wiki)
7. Gizmo (more info at http://responsphere.calit2.net/gizmo)
8. UrbanMesh
9. ESB for System Integration (more info at http://www.dsscc.com/)
10. LBT: Location Based Tracking
11. P2P Notification (more info at http://traffic.calit2.net and http://traffic.calit2.net/communicator)

MetaSim

Metasim is the result of integration of diverse simulation capabilities being developed at different RESCUE partners into a single integrated system. In particular, Metasim is an amalgamation of the transportation simulation with micro-level agent simulator (DrillSim being developed at UCI) and a cellular infrastructure simulator being developed at UCSD. Metasim provides RESCUE researchers mechanism to test and validate IT solutions in a very rich set of scenarios which neither of the three simulators could provide individually.

Activities and Findings.

Many project RESCUE efforts include a simulation component, and integration of the various simulators has a potential for synergy. Given this opportunity, the transportation test bed platform was expanded in scope to encompass existing simulation efforts at multiple research sites and became project MetaSIM. MetaSIM was a natural extension of the ongoing collaboration between universities and researchers forged by Project RESCUE, and the transformation of the transportation testbed intoMetaSIM formalized efforts to address issues surrounding model integration. This section of the annual project report describes the purpose of the MetaSIM testbed, and the current status.

METASIM is envisioned as 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 as measured by reduction in expected losses, evacuation times, and other impacts can be determined for each technology developed. Outside of the research community, METASIM will prove useful to emergency managers and first responders by providing centralized and wireless dissemination of disaster simulation data and information. Before an event, disaster simulations of probable events will aid in the prioritization of mitigation activities and increase preparedness through training scenarios. Immediately after an event, METASIM will aid in situational awareness and resource deployment. During the recovery phase, METASIM will help assess long-term shelter and public assistance requirements. METASIM will be a highly distributed collection of simulation models connected by a series of translators that access a common database. Simulation models running on computers at different locations are connected together to create a meta-simulation capability that allows the assessment of the larger impacts of a disaster. Models simulating various aspects of a complex event do not have to coordinate data transfer with each other, but are able to share simulation results through a common enterprise database. A standardized communication protocol is envisioned to allow other simulation models to be incorporated without system redesign. Through the web interface, a basic user would be able to run each module as needed, whereas researchers could incorporate meta-simulations into their own projects by calling individual models as a computer service. By utilizing this distributed simulation environment, METASIM has the potential to become a comprehensive modeling platform for simulating the effects of large disasters.

End users of the envisaged system are: 1) Earthquake engineers and scientists, sociologists, geographers, computer scientists, and other domain experts; 2) Civil, structural and

transportation engineers; and 3) Emergency management professionals. Anticipated applications include: 1) Testing emerging technologies; 2) Training based on probable scenarios; 3) Drills from probable scenarios, such as those conducted by EOCs;
4) Post-event loss estimation and impact analysis; 5) Post-event simulation of cascading events; 6) “Plug and play” integration of additional simulation tools; and 7) Assessing mitigation strategies and alternatives.

Significant progress has been made towards developing a roadmap for future model integration. A guiding principle of MetaSIM has been that if the various simulation modules developed in the RESCUE project could share data in real time, they would become more than the sum of their parts, and that a platform and protocol supporting modular and extensible integration. The following list provides a summary of key principles that will guide the MetaSIM testbed. They represent a common understanding amongst researchers and developers on overriding principles for future work on MetaSIM. Adherence to these principles will be key for project success.

a) Modularity and extensibility- MetaSIM will serve as a common database for visualizing the results of the various modules and as a communication hub, storing and facilitating the translation of data between individual components. This communication portal will not serve as a centralized database for all the individual models. Each model is likely to have specific data requirements that may or may not correspond with other models. The scalability of MetaSIM will depend on each individual model sharing a reasonable amount of data through the centralized system.
Ultimately, MetaSIM will be a collection of plug-and-play simulation tools connected by a series of translators and a database. With proper definition of inputs, outputs, timing, and scale, the results of each simulation component could 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 require tight integration.

b) Integration of key components- MetaSIM will consist of a suite of simulation tools. Given that there are likely to be many overlapping features in the various simulators, such as an interface, a method to view results, and a database, many of the components of a given simulator will not be used, and advanced users may want to use the various simulation engines outside of MetaSIM.

c) Analysis at multiple scales- The various simulation tools will integrate results from micro-simulations of very small areas to large-scale statistical models covering very wide regions. A key component of MetaSIM will be the integration of these various levels of analysis, so that micro-scale benefits are extrapolated to regional effects, and regional effects are used to inform micro-simulations.

d) Simplified user interfaces- Each simulation component of MetaSIM draws from many disciplines, and expert use requires extensive study. However, the vision of MetaSIM is of a product that can be used with very little, if any, expertise in the science or technology that supports MetaSIM. This will be accomplished through extensive use of defaults, so that users can adjust a minimal number of parameters of interest without dedication of significant resources. Table 1 illustrates how this goal is accomplished through a series of “user levels”.

Level	Who	Time Required	Input	Results
L1	Emergency Manager	Minutes	Whether or not to run various modules	Default output from each model at some level of aggregation, with IMS or other interface to detailed results
L2	Researcher Response Personnel	Hours	Select from various predefined options, representing a baseline, and the integration of technology.	Comparison of various outputs with and without integration of technology.
L3	Advanced Research	Days	Define custom input, integrating a new suite of technology.	Comparison of various outputs with and without integration of technology with new technologies analyzed.
L4	RESCUE Researchers	Weeks	New regions, models, or geography.	Comparison of various outputs with and without integration of technology, analyzing a new area.

Preliminary identification of anticipated users of DrillSim within MetaSIM, indicating the expected level of effort and data requirements to complete an analysis.

For MetaSIM to achieve the ultimate goals of modularity and extensibility, many integration issues must be resolved. MetaSIM will address these issues through leveraging existing resources within the RESCUE project, and focusing on a key application: modeling the benefits of integrating cellular technologies during evacuation. This application will address four simulation modules currently being developed and/or used within the RESCUE project. The crisis simulator (Inlet) will provide initial estimates of damage throughout a region. Based on the estimated damage to buildings and the cellular infrastructure, DrillSim will model evacuation of the UCI campus. This evacuation will consider a default occupancy level for the campus, for every floor of every building, and will be modeled both with and without IT technologies or hardening of the cellular infrastructure. As evacuation occurs at a building level, evacuees will enter an out-door campus level evacuation, from which they will continue to an automobile-based evacuation using the transportation model. A cellular simulation tool will consider how to optimize the remaining cellular load to facilitate evacuation. The vision of this initial deployment project has evolved amongst project team members, given the capabilities and limitations of existing models. A common, application-based focus is emerging that will provide a concrete deployment challenge to address issues of modularity and extensibility.

Products.

1. Synthetic Humans in Emergency Response Drills Daniel Massaguer; Vidhya Balasubramanian; Sharad Mehrotra; Nalini Venkatasubramanian; 2006 AAMAS Conference, Demo Paper, 2006.
2. Multi-Agent Simulation of Disaster Response Daniel Massaguer; Vidhya Balasubramanian; Sharad Mehrotra; Nalini Venkatasubramanian; ATDM Workshop in AAMAS 2006.
3. DrillSim: A Simulation Framework for Emergency Response Drills Vidhya Balasubramanian; Daniel Massaguer; Sharad Mehrotra; Nalini Venkatasubramanian; ISI 2006.

4. Development of a Web-based Transportation Modeling Platform for Emergency Response. Sungbin Cho; Charles Huyck; Shubharoop Ghosh; Ronald Eguchi; 8th U.S. National Conference on Earthquake Engineering, San Francisco, 2006.

Contributions.

The Transportation Simulator

The transportation simulator provides a platform for testing the assumption that rapid and accurate information conveyed to drivers can significantly improve the performance of a transportation system when viewed within the context of evacuation. Drivers weigh information reliability for: 1) understanding the status of an event; 2) assessing the expected duration of disruption; and most importantly; 3) determining the best action to take to minimize the impact. With additional and more accurate information, it is expected that drivers will behave in a more predictable manner and generally follow instructions and directions from ITS (Intelligent Transportation Systems), and thus will be able to reduce average travel times during an emergency. The transportation module consists of an integrated model of simplified quasi-dynamic traffic assignments, and a destination choice model. In the last year, information that will become available through IT solutions is synthesized through parameters, such as information reliability, rapidness of dissemination, penetration rate and degree of customization, to reduce uncertainties associated with decision making when evacuating a congested network.

In the current proof-of-concept, three parameters characterize the type of information that might be provided to drivers. The *reliability* of information indicates “how much” the drivers trust the warning given to them for evacuation. The more reliable the information, the more people depart sooner. The *timing* identifies the lag between the disaster and the notice to evacuate. The *customization* refers to the degree of information specificity with regard to their current geographic location. If no customization is applied to the evacuation notice, people in a given block will all receive the same messages. However, if an IT solution can detect the location of a cellular phone, and it may be possible to give individuals very specific routing information.

The trip generation model allocates population within the exposed area over time, depending on *Reliability* and *Timing*. In the current model, network congestion is not considered in trip generation. Given *Reliability*, r, and *Timing*, T, trip generation at any given simulation time period t, g_t is calculated by Equations 1 and 2. By the warning given at time T, the population moves at its maximum rate g_{max}. Equation 1 specifies a linear relationship between reliability and the maximum generation rate. The model assumes a gradual increase in the number of evacuees over time, according to Equation 2. The key parameters are available as adjustable inputs to the model, for users to assess the efficacy of different methods of integrating IT into emergency response. The module has been tested with various small scale evacuation scenarios.

$$g_{\max} = \alpha + \beta \cdot r \quad (1)$$

$$g_t = \begin{cases} g_{\max} \cdot t/T, & if\ t < T \\ g_{\max}, & if\ t \geq T \end{cases} \quad (2)$$

The transportation model used to assess route choice for evacuation consists of a series of database queries. The basic algorithm implemented is detailed below. Drivers continuously depart from an initial location (Step 2), and take the best route at that instance (Steps 5 and 6). On each route, traffic is assumed to be continuous with uniform density. Step 7 counts how many drivers are exiting the evacuation area within each time period. Residual evacuation demand is added to the trip generation module for the next time period, and is factored into the congestion calculations in Step 3. The simulation is repeated for a predefined number of time periods.

Step 0: Initialize: Clear storages for intermediate calculations
Step 1: Estimate trip generation in each time period
Step 2: Increase simulation time period by one
Step 3: Update congested link travel time based on traffic volume on the link
Step 4: Calculate path travel time between all zone-pairs, by aggregating link travel times in each route
Step 5: Choose best travel route-destination from each origin to destination-route combination according to the level of customization.
Step 6: Assign the trips to routes based on the allocated demand to the links in the selected path
Step 7: Calculate evacuated trips and evacuation demand in the system
Step 8: Evaluate stopping criterion. Stop analysis if current time period reaches the simulation duration

Resolving changes in driving behavior and IT solutions is still problematic. A premise for the transportation testbed is that IT solutions may improve disaster response. By providing means for rapid assessment of the situation and optimal plans, online tools should improve emergency responses extensively, but the link between IT and behavior is currently left to conjecture. This will be addressed in year 4 through integration with MetaSIM. Additionally, for better representation of drivers' mobility, the transportation module will be modified to incorporate meso or micro scale traffic dynamics.

DrillSim Simulator

DrillSim is designed as a multi-agent crisis response activity simulator which has plug and play capabilities. The goal of DrillSim is to test IT solutions in the context of disaster response. The activity is modeled at an individual level, and software agents model humans. A neural network based decision making system models the human decision making in our system. DrillSim version 1.0 which simulates an evacuation within a building has been developed, and DrillSim version 2.0 which simulates an evacuation at a campus level is being implemented currently. Interfaces to input data for the DrillSim and a visualization interface have been developed. Virtual reality-augmented reality integration is being pursued currently, and an initial level integration has been achieved by sending the output of DrillSim to a real person carrying a PDA or a mobile computer.

In the past year, models for representation of space and the impact of different phenomenon on space have been studied. The models are at different resolutions to capture accurately the impact and also keep the system scalable. Planning over multiple resolutions of spatial data, and efficient data management techniques exploiting the multiple resolution representation of space is currently being explored.

The software agents, which model human beings in DrillSim, are the information processing entities. Each agent has a view of the world, based on which it makes decisions, and plans to execute them. The decision making process is modeled using a neural network. The behavior of agents form the basis of the activity modeling, and it provides a flexible way to implement different activities by just modifying behavior. Agents also assume different roles, and the behavior is dependent on the roles. DrillSim has the facility to edit existing roles and also add new roles into the system.

Additional progress includes the creation of base data sets, which are being used by DrillSim and other projects within Rescue. Maps and GIS data for the entire UCI campus, and detailed building CAD maps have been analyzed and stored in the database. These datasets support a geographical model for DrillSim. Data of drills (video of drill, questionnaire etc) conducted in UCI have also been stored. DrillSim activity modeling is being calibrated using data from these drills.

Designing agents that realistically model human beings is the primary challenge in DrillSim. In general, humans use information and knowledge to take decisions, which are rarely simple and independent from decisions made by others. Additionally, when faced with identical information, a group of individuals may decide to react differently, based on knowledge. Factors like social network, risk adversity, resources, and how people react to technology impact the decision making process. Work is being done to improve the behavior of agents in DrillSim.

DrillSim models activity at the individual level, and since every agent makes a decision at every instance that impacts the output of the simulation, there is a significant scalability issue. Research into scalability solutions has included modifying the scale of resolution of data and interspersing macro level modeling with micro level activity modeling. In year 4, DrillSim version 2.0 will be deployed for campus level evacuation, and integrated into actual CAMUS drills. Interfaces for DrillSim will be designed and integrated, and the simulation module will be integrated with MetaSIM

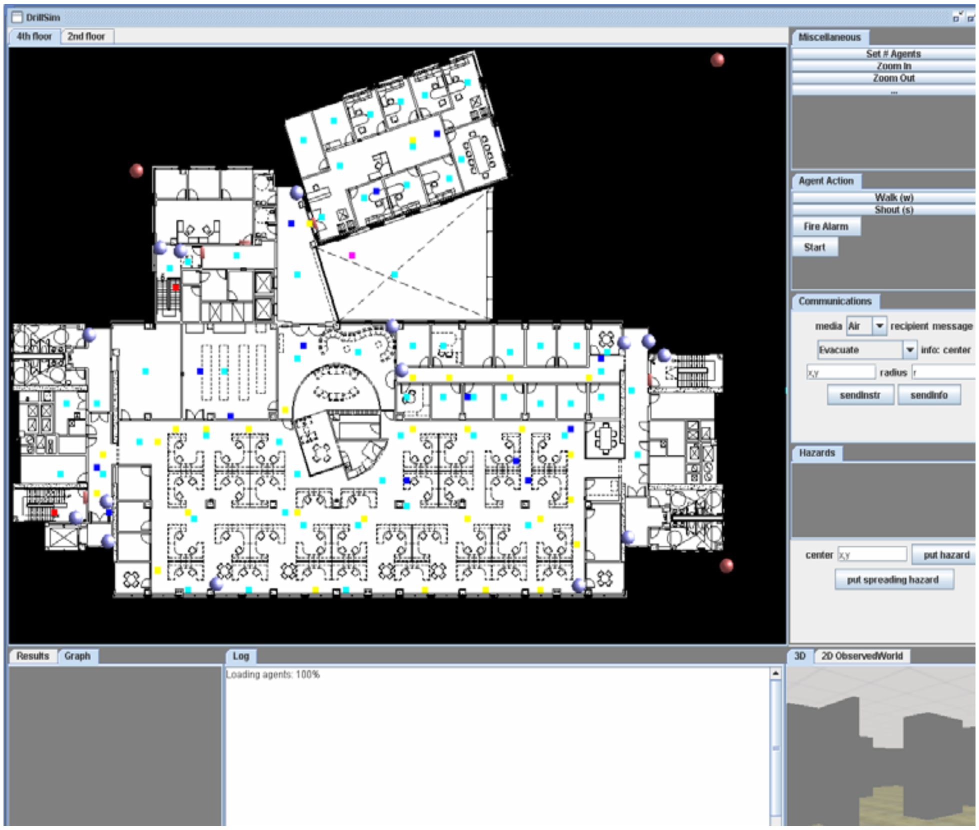

DrillSim

Crisis Simulator (InLET)

In the years following the 1994 Northridge earthquake, many researchers in the earthquake community focused on the development of loss estimation tools, such as HAZUS (HAZards-US). Advances in GIS technology, including those associated with desktop GIS programs and robust application development languages enabled rapid scenario generation that previously required intensive data processing work (in many cases, manual activities.) In some cases, these programs were linked with real-time earthquake information systems, such as the Caltech USGS Broadcast–of-Earthquakes (CUBE) system, or ShakeMaps (USGS) to generate earthquake losses in real-time. The Federal Emergency Management Agency (FEMA) funded the development of a high-profile and widely-distributed loss estimation program called HAZUS (HAZards-US). Because of its flexibility in allowing users to import into the program custom datasets, the use of HAZUS in real events (e.g., 2001 Nisqually, Washington earthquake) sometimes resulted in conflicting results, especially when different users were involved in modeling the same event.

This issue (i.e., conflicting results) highlighted the need for a centralized system where data and model updates could cascade to multiple users. In the past year, a successful implementation of the Internet-Based Loss Estimation Tool (InLET) was developed. InLET is the first known web-based massive simulation program that incorporates: 1) an earthquake-based disaster simulation module; 2) building and infrastructure damage estimation routines; 3) socio-economic loss estimation models, and 4) a transportation routing and evacuation module. Currently, the model works on a database for Los Angeles and Orange Counties. The program has a mapping interface with standard GIS functionality such as re-center, and zooming in and out. Additionally, the basic framework easily supports overlaying various vector and raster data, including points of interest, data gleaned through online searches, or satellite imagery and aerial photography. Development has progressed as planned, and the product has been demonstrated for use by multiple end users at one time.

A primary challenge has been structuring large database files used for computation so that the results are usable by multiple users over the web concurrently. Additionally, allowing multiple users to access the system for extensive analysis makes the scalability of InLET an issue. The models are mainly based on computationally intensive SQL queries. Over a series of improvements, InLET is being migrated to a robust analytic system that is based on dependable SQL Server. The queries are being formatted to simultaneously work for multiple users. Primary functionality is being built to incorporate USGS ShakeMaps and ShakeCast results.

InLET is designed to be used by first responders, planners, and anyone involved with emergency response. It is a tool to be used for someone to see where the damage will be likely to occur and how one should plan accordingly. Outreach has included GIS analysts at city, county, state, and federal emergency management agencies. In year 4, InLET will be modified to accommodate multiple users simultaneously in a simulation capacity. The backend database will be migrated to SQL Server to provide a more robust architecture. INLET Online loss estimation tool: http://rescue-ibm.calit2.uci.edu/inlet/

Additional Responsphere Papers and Publications

The following list of papers and publications represent a list of additional 2005-2006 research efforts utilizing the Responsphere research infrastructure:

1. Emotion Detection from Speech Signals and its Applications in Supporting Enhanced QoS in Emergency Response Rajesh Hegde; Manoj B.S; Bhaskar Rao; Ramesh Rao; ISCRAM 2006, 2006.
2. Exact Bounds for Degree Centralization. Carter Butts; Social Networks, 2006.
3. Iterative Joint Source-Channel Decoding of Speech Spectrum Parameters over an Additive White Gaussian Noise Channel. A. D. Subramaniam; B. D. Rao; W. R. W. R.; IEEE Transactions on Speech and Audio Processing, 2005.
4. l0-norm Minimization for Basis Selection. D. Wipf; B.D. Rao; Neural Information Processing Systems (NIPS), Vol. 17, 2005.
5. Recognizing Cars. Louka Dlagnekov; Serge Belongie; Technical Report, 2005.
6. Geopolymer based Sensing. Peter Chang. NSF Workshop for Earthquake Mitigation. Ankara, Turkey. 09/2005.

Courses

The following undergraduate and graduate courses are facilitated by the Responsphere Infrastructure, use Responsphere equipment for research purposes, or are taught using Responsphere equipment:

UCI ICS 299, Loud and Clear Project, Tsudik, Summer: 2005.
UCI ICS 280, Secure Group Communication, Tsudik, Spring: 2005.
UCI ICS 280: System Artifacts Geared Towards First Responders, Mehrotra: Spring 2005.
UCI ICS 214A, Principles of Data Management, Mehrotra, Fall:2005.
UCI ICS 214B, Distributed Data Management, Mehrotra, Winter: 2006.
UCI ICS 290, Research Seminar, Mehrotra, Winter, 2006.
UCI ICS 215, Advanced Topics in Data Management, Mehrotra, Spring: 2006.
UCI ICS 203A, Introduction to Ubiquitous Computing, Lopes, Winter: 2006.
UCI ICS 278, Data Mining, Smyth, Spring: 2006.
UCI ICS 199, Directed Research, Venkatasubramanian, Fall, 2005.
UCI ICS 280, Systems Support for Sensor Networks, Venkatasubramanian, Winter: 2005
UCI ICS 290, Research Seminar, Venkatasubramanian, Winter: 2006.
BYU, Access Control in Open Systems, Seamons, Winter: 2005.
UCSD ECE 191, A Robust Multi Modal Speech Recognition System, Rao, Winter 2006.
UCSD ECE 191, An Embedded Speech Recognition System, Rao, Spring 2006.
UCSD ECE 191, Designing a High Capacity Wireless Mesh Network, Manoj, Spring 2006.
UCSD ECE 291, Zigzag Tactile Smart Pointer Guidance System for First Responders, Miller, Spring 2006.
UCSD ECE 291, Engineering Group Design Project, Hegde, Spring, 2006.

UCSD MAE 156, Mesh Network Antenna Caddy, Kimball, Spring, 2006.
UCSD ECE 158, Data Networks II, Manoj, Winter, 2006.

Equipment

The following table summarizes the types of equipment the UCI Responsphere team obtained for the project. The most significant purchases in year 2 were the 8 processor Operton server (with AMD's generous assistance), additional video sensors, the EvacPack system, the outdoor Phase 1 wireless infrastructure, and the expansion of the people counting system., The Opteron server was obtained by Responsphere for $22k. An identical machine (with AMD dual core processors) retailed from Sun Microsystems for around $70k. Additionally, a partnership with CDWG (retailers for Canon and Linksys) resulted in significant cost savings for cameras pricing. In all cases, education pricing and discounts were pursued during the purchasing process.

	UCI Equipment	
Date	Equipment	Usage
9/1/2005	Linksys WVC54G	Smart-Space Cameras
9/1/2005	Linksys WAP-POE	POE Kits: Smart-Space Cameras
9/21/2005	Appro	Opteron Server - Simulation Server
9/21/2005	Appro	Opteron Server - On-Site Maintenace
9/21/2005	APC	UPS for server
9/20/2005	LDC	TDT4 Data Set
10/1/2005	HomeSecurityStore	People Counters
10/1/2005	Walker Wireless	People Counters
10/1/2005	Paralax	People Counters
10/1/2005	MicroOptical	VR/AR visualization
10/1/2005	Cappuccino PC	EvacPack Computing
11/1/2005	Fry's Electronics	EvacPack Computing
12/1/2005	V-Realites	EvacPack Computing
12/1/2005	SoftTouch Inc.	Closed-Captioning Text Extraction
12/1/2005	CDWG	Closed-Captioning Text Extraction
1/1/2006	NACS	Co-Location for network Backup Storage - Yearly
1/1/2006	CDWG	Buffalo Terrastations
1/1/2006	CDWG	Disk Drives for off-site data storage
3/1/2006	5G Wireless	IEEE 802.11 Outdoor Wireless for Smart-Space
3/1/2006	PowerStream	power supplies for mobile cameras
3/15/2006	Sparton	Avionics: EvacPack navigation
3/15/2006	CDWG	Smart-Space mobile cameras
3/28/2006	Draeger	Gas Monitors for Smart-Space
4/2/2006	Dell	Grad. Student PCs
4/7/2006	CDWG	Outdoor cameras for Drills
5/8/2006	Point Grey	Firewire Cameras
5/10/2006	PowerStream	Power supplies
6/18/2006	Cobra	First Responder FRS radio systems
6/20/2006	Echelon	I-LON PL Server

At UCSD, the nature of equipment purchased is very different, as the majority of the equipment is being fabricated from components by the UCSD team. Approximate anticipated year 2 expenditures (plus applicable overhead) are as follows: $50K on supplies used to fabricate CalMesh nodes; $80K on components for the Portable Tile Display Wall; $1K on supplies for the vehicle; $15K on supplies used to fabricate a permanent mesh network infrastructure for the GLQ testbed; $7K on RF and GIS modeling tools and equipment; $60K on generic project supplies and equipment; $65K on staff and undergraduate salaries to support the fabrication of this equipment; $3K on travel to send project personnel to relevant conferences and symposia, and meetings with collaborators at other institutions.

UCSD Equipment Fabrications		
Quantity	Equipment	Usage
32	CalMesh Nodes	Custom wireless nodes fabricated by UCSD ResponSphere team.
4	Gizmo Nodes	Test platform for self-deploying, remotely controllable wireless network nodes equipped with various kinds of field sensors
n/a	Vehicle Components	Components purchased and fabricated include: Antenna Caddy WiFi Bullhorn Satellite Dish & Deployment System
1	Portable Tile Display Wall (fabrication)	Components include: 14 bright-screen tablet PC laptops 2 bright-screen, weatherproof HDTVs 4 clustered server PCs to drive display Mounting hardware
3	RF Modeling Tools	OpNet TEMS Spectrum Analyser
2	PCs	General use workstation/server and laptop