



Secure-CITI Critical Information-Technology Infrastructure

Daniel Mossé, Louise Comfort*, Ahmed Amer, José C. Brustoloni, Panos K. Chrysanthis,

Milos Hauskrecht, Alexandros Labrinidis, Rami Melhem and Kirk Pruhs

University of Pittsburgh

Computer Science Department and *Graduate School of Public and International Affairs

The Secure and robust Critical Information-Technology Infrastructure (S-CITI for short) project aims at providing support to Emergency Managers (EMs) that are faced with management of resources and with decisions before, during, and after emergencies or disasters. Our approach consists of using new and existing sensors to gather data from the field, processing this data to detect and predict emergency/disaster situations, and disseminating this data among the appropriate organizational units. The data flow will be done in a reliable and secure manner and EMs will coordinate actions in a *Virtual Coordination Center* (VCC), which need not be in a fixed (and thus vulnerable) physical location. The EMs are responsible for indicating what type of data is more valuable, so that S-CITI can display that information appropriately.

There are nine faculty members involved in this effort and five research groups, with several interconnected areas of research. In this paper, we present the "big picture" and then give a brief overview of the current subprojects and results.

The current primary mode of operation in disaster detection and response is through the 9-1-1 system, where humans call in emergencies and the appropriate personnel is dispatched. This detection/response system has proved adequate in many situations, but slow in other scenarios (especially when the humans themselves are involved in the emergency/disaster). Further, counting on reports from the field may be inadequate when communication breakdown exists between the area affected and "the other side" or among personnel responding to the emergency/disaster. A good example of such communication and infrastructure breakdown is the unfortunate levee breaks that happened in New Orleans in 2005 [1].

Our system supports, through information technology, the coordination of usage of existing resources and distribution of the data to different organizational units. The VCC facilitates efficient and quickly coordinated actions to natural and human-caused disasters [2]. The data collected and the actions taken are analyzed through a learning module that will feed post-emergency data into a pre-emergency decision-making module (improving response in the next emergency).

- Supported in part through NSF-ITR award ANI-0325353.

The *socio-technical research team (STRT)* combines expertise from data management to social networks, because the issue of Emergency Management is as much a technical challenge as it is a social networks problem. One of the main strengths in our group is the ability to create a connection between the technical and the Emergency Personnel (especially in the Pittsburgh area). Toward the VCC goal, we have developed the *ISIS Executive Dashboard* [3] that provides real-time decision support to practicing EMs during disaster situations.

The STRT developed a preliminary version of a patient-tracking module (see companion powerpoint file), which was demonstrated in a simulated operations exercise planned by the Counter-terrorism Task Force of Southwestern Pennsylvania. This module tracks the information flow from a triage site to a transportation officer to the Receiving Room of an Emergency Medical Department in a hospital. The module supports the exchange of information at all three decision points in a multi-way communication process to enable the timely transport of patients to medical care in the most efficient, informed manner.

Based on this experience and critical evaluation of the STRT, the patient tracking module was revised to include communications among multiple sites of decision making for tracking patients. In essence, the system must enter data and allow monitoring applications to continuously mine data streams for interesting, significant, or anomalous events (e.g., patient data being entered in the system, hospital data being updated, etc). The data must be maintained fresh (not stale data) and the queries must return the most up-to-date values possible. Mechanisms to explore this situation have been reported in [4, 5, 6].

The *ad-hoc network research group (ANReG)* focuses on (a) development of network protocols for energy-efficient data access, (b) accessibility to data even in the presence of disconnected networks, and (c) security of the network.

Because power is a mainstay of ad-hoc mobile networks, which are and will be the type of devices/networks used by emergency personnel (cell phones, PDAs, etc), our group has been investigating networking algorithms (MAC and routing) to minimize the energy consumption in these networks [7]. Our adaptive algorithms have demonstrated significant potential at increasing the longevity of networks used by EMs, relying on optimizing the inquiries to which field officers and volunteers are attempting to respond (e.g., how many beds there are in the ICU of hospital H, what are the traffic conditions to the hospital, etc). Prediction algorithms also contribute to the decrease of the energy and power consumption in these networks, which is a subject of our ongoing investigation. Lastly, our resource management algorithms allow each EM to specify for him/herself the importance of each datum, so that the system can supply the different EMs with the most appropriate data in the most efficient

way. In particular, we will dynamically synthesize databases with new consistency and authentication procedures bound to the type, size, and importance of the data.

The second issue is a new trend in network protocols that abandons the assumption that the network is always connected. We study how couriers (mobile message-forwarding nodes) can enable communication in partitioned networks, carrying messages physically between partitions. Our contribution is a new, cross-layer routing approach based on the observation that orders from an EM (leader sending tasks to responders) also controls responders' mobility and thus their ability to forward messages. We schedule responders' movement considering both their tasks and network needs. Our simulations demonstrate performance benefits of our approach in a variety of scenarios [8]. We have also developed a prototype EOC-like testbed with sensors and mobile nodes for experimentation in this area.

The third issue is that, in mobile devices, not only energy and connectivity are important factors, but the issue of security is especially important because the network infrastructure may be under attack. For that, we are investigating two problems. First, when couriers that receive messages from or send messages to the EOC are being jammed. We have modeled this scenario through a Markov Model, in which the state of the couriers can be blocked, free to receive, or moving (to another location to avoid jamming). Our simulation and modeling results allow for determining the exact break-down point, that is, when the EOC should either add couriers, or suffer from lack of connectivity [9]. Our second aim is to provide security at the level of the network, not allowing messages to be forged, intercepted, deleted, or inserted by attackers. Clearly, our protocols respect the different organizational units and the privacy of data.

The *intelligent monitoring and diagnostics group (IMDiG)* has focused on modeling static and dynamic dependencies in large distributed systems with continuous and discrete random variables, and efficient optimization of control activities in such distributed systems. Because transportation systems are one of the key elements of the emergency response system, one of the goals of our work is to build statistical models that let us represent spatial correlation patterns among traffic variables (speed, volume, throughput, etc) over the network under different conditions (free-flowing normal, traffic congestion). Such models support EMs under a variety of distress conditions. For example, understanding congestion patterns and traffic predictions would allow the EMs to better route/dispatch emergency resources (ambulance, fire, repair units) in a way that reduces response time.

We have obtained two years worth of traffic data from over 100 sensors placed on major Pittsburgh highways and have investigated statistical properties of the data and relations among sensors themselves. We have built initial multivariate models that can lead to, for example, prediction of different events. A proof-of-concept algorithm was developed that successfully detects anomalous days from vehicular count data. Our concrete goal is to develop an online system that predicts with high probability a possible accident before it happens, by simply using inferences on the data obtained from the sensor network in place.

Finally, the *advanced data management technologies laboratory (ADMT)* in collaboration with the *storage research group (SRG)* have produced a wide range of data management

algorithms, focusing in quality of data (QoD) and quality of service (QoS), in the presence of resource constraints, which characterize emergency response environments. For example, at the point of generation of data (e.g., from networks of sensors or mobile devices), we have proposed energy-efficient data acquisition techniques [10], which combine in-network processing [11] and in-network, data-centric storage [12]. To effectively propagate data within such sensor networks, we proposed semantic-based, multi-criteria, self-optimizing algorithms for constructing efficient and robust routing topologies [13]. Data in the emergency management domain typically come in the form of data streams that need to be continuously monitored for interesting or anomalous events. Towards this, we have proposed QoS- and QoD-aware techniques for processing of continuous queries [4, 5] and implemented some of these ideas in our prototype Patient Tracking module, mentioned above. The produced data is disseminated to mobile users, such as first responders, who utilize our proposed energy-efficient mobile caching and energy-safe prefetching algorithms, which are self-optimizing for changing workloads [14].

REFERENCES

- [1] Comfort, L.K. *Communication, Coherence, and Collective Action*. Public Works Management & Policy., Sep. 2006.
- [2] Berfield, A., Chrysanthis, P.K., Labrinidis, A.. *Automated Service Integration for Crisis Management*, 1st Int'l ACM Workshop on Databases in Virtual Organizations, 2004.
- [3] IISIS prototype, <http://www.iisis.pitt.edu>
- [4] Sharaf, M., Chrysanthis, P.K., and Labrinidis, A. *Preemptive Rate-based Operator Scheduling in a Data Stream Management System*. IEEE AICCSA, 2005.
- [5] Sharaf, M., Labrinidis A., Chrysanthis P.K., and Pruhs, K. *Freshness-Aware Scheduling of Continuous Queries in the Dynamic Web*. ACM WebDB, 2005.
- [6] Qu, H., Labrinidis, A., and Mossé, D. *UNIT: User-centric Transaction Management in Web-Database Systems*. IEEE ICDE 2006.
- [7] Gobriel S., Melhem, R., and Mossé, D. *BLAM: An Energy-Aware MAC Layer Enhancement for Wireless Adhoc Networks*. IEEE WCNC, 2005.
- [8] Brustoloni, J., Khattab, S., Santamaria, C., Smyth, B. and Mossé, D. *Integration of Application-Layer Scheduling and Routing in Delay-Tolerant MANETs*. Pitt CSD T.R., 2006
- [9] Khattab, S., Mossé, D., Melhem, R. *Honeybees: Combining Replication and Evasion for Mitigating Basestation Jamming in Sensor Networks*. WPDRTS, 2006.
- [10] Sharaf M., Beaver J., Labrinidis A., Chrysanthis, P.K. *Balancing Energy Efficiency and Quality of Aggregate Data in Sensor Networks*, The VLDB Journal, Dec. 2004.
- [11] Xia, P., Chrysanthis, P., Labrinidis, A. *Similarity-Aware Query Processing in Sensor Networks*. WPDRTS, 2006.
- [12] Aly, M., Morsillo, N., Chrysanthis, P.K., and Pruhs, K.. *Zone Sharing: A HotSpots Decomposition Scheme for DataCentric Storage in Sensor Networks*, DMSN, 2005.
- [13] Li Q., Beaver J., Amer A., Chrysanthis P., Labrinidis A., Santhanakrishnan G. *Multi-Criteria Routing in Wireless Sensor-Based Pervasive Environments*, Journal of Pervasive Computing and Communications, Dec. 2005.
- [14] Larkby-Lahet J., Santhanakrishnan G., Amer A., Chrysanthis, P.K. *STEP: Self-Tuning Energy-safe Predictors*, ACM/IEEE MDM Conference, 2005.