

# Crisis Management in a Data Fusion Synthetic Task Environment

**Peter D. Scott**

Dept. Computer Science and Engineering  
University at Buffalo  
Buffalo NY 14260-2000  
USA  
peter@buffalo.edu

**Galina L. Rogova**

Encompass Consulting  
9 Country Meadows Drive  
Honeoye Falls, NY 14472  
USA  
rogova@rr.rochester.com

**Abstract** - *The core problem in the decision sciences has always been how to use scarce data optimally. But the proliferation of inexpensive computer-based instrumentation and broadband communication make it commonplace in current applications that data is abundant, often excessive. The technical challenges are twofold: how to distill data into knowledge, and how to use that knowledge wisely. The former is the problem of information fusion, and this communication reports progress in design of a fusion system for the data-rich knowledge-poor environment following a natural disaster, namely the initial response phase after an event similar to the Northridge CA 1994 earthquake. Use of a synthetic task environment permits a high degree of control and evaluation, and allows emphasis on fusion methodology. Higher level fusion products are derived from user needs, based in part on the California Standard Emergency Management System framework using cognitive work analysis and positioned within a disaster response ontology.*

**Keywords:** Data fusion, synthetic task environment, higher level fusion, disaster management, earthquake, ontology, cognitive work analysis.

## 1 Introduction

When disaster strikes, the first reflex of both victims and response teams is to draw inward, act in ways that are most direct and familiar. Civilians in the disaster area rely first on family and neighbors, emergency responders on their partners and immediate unit organizations. Extensive information sharing or other cooperation between separated units is seldom noted in the early disaster response stage, becoming more common only in the follow on recovery and reconstruction phases.

The potential benefits of wider cooperation have historically been mitigated by the unavailability of robust communications or transportation systems to support them. Increasingly, however, this barrier is being eliminated. More complex and potentially superior disaster response strategies become feasible as advances in these technologies empower fast, widespread information and resource sharing within the disaster zone and between that zone and neighboring areas.

In particular, rapid dissemination and exploitation of usable information can be of great benefit in the first few hours after disaster strikes. Accurate wide-area damage, casualty and resource assessments can be used to direct responders effectively, guide evacuation efforts and position assets.

A key enabling technology for effective information exploitation is information fusion. Reports will begin flowing from observers, responders and instruments of various descriptions that must be parsed effectively and integrated quickly. In the "fog of disaster," reports will be of varying significance and accuracy, perhaps redundant or contradictory. Reporting protocols, as well as ontological interpretations, will differ among units and jurisdictions.

This represents both a significant opportunity and a real challenge to the data fusion community. To focus effort on the most relevant issues confronting data fusion research, the AFOSR program in Information Fusion has identified three major challenge problems. In a FUSION 2002 paper, Tangney [1] delineated these three problems. We have set a response to the Crisis Management Challenge Problem in the context of natural or man-made disasters, and in particular to explore the design and evaluation of data fusion techniques in the immediate post-earthquake disaster period within a synthetic task environment. This communication reports progress on our effort since the previous report by Llinas [2].

### 1.1 General approach

In determining our approach to the Crisis Management Challenge Problem, tradeoffs had to be considered. Domain specificity vs. reusability was one such tradeoff. Emphasis on the former would require immersion in a narrow problem environment, the creation and deployment of object models, dynamic processes and human-machine interfaces with rigorous "custom" domain fidelity. The latter would emphasize methodological implications: abstraction and reuse of the specific challenge problem solution architectures and algorithms in other higher-level data fusion applications, and the migration paths by which such reuse could be achieved. Our choice was to emphasize reusability by setting modeling fidelity at a level which avoids over fitting the solution to narrow phenomenological details, while at the same time maintaining its basic plausibility.

A related consideration was the selection of a specific crisis scenario. Given the desired reusability emphasis, it was required to be representative of a broader applications context. In addition, emphasis on methodology rather than

domain specifics implied that we would rely on available models and databases. In the end the context of natural and man-made disasters was selected. This has direct relevance to both civilian and military domains, including Department of Homeland Security applications and the Joint Battlespace Infosphere [3] fusion environment. More specifically, an early phase earthquake response scenario was selected.

Another early design tradeoff was the choice of a real, synthetic or hybrid environment. Each has its advantages, but in the end we judged the synthetic environment to best fit our other goals and emphases. The benefits and challenges of synthetic earthquake environments have been discussed by Takeuchi [4]. As noted by Llinas [2], the high degree of experimental control in a synthetic environment promotes consideration and controlled testing of performance and of performance metrics. Nomination, algorithmic implementation and testing of measures of performance and effectiveness are priority problems in the design of high-level data fusion systems [5], and the choice of a synthetic task environment supports this line of inquiry.

There are several relatively distinct phases to the post-earthquake response. The initial stage is dominated by rescue and medical management of existing casualties, together with efforts to mitigate the risk of suffering additional casualties. We refer to this as the First Response Phase (FRP). Subsequent phases focus on supplying food and shelter to those displaced by the earthquake, survey and repair of general infrastructural damage, protection of damaged property, restoration of normal public services, and restoration of normal economic activity. Each phase has its own key entities, events and processes. In order to simplify domain modeling, we chose to consider only the FRP, which we take to extend no more than 24 hours after the shock which initiates it.

Earthquakes consist of a primary shock event followed by a series of aftershocks occurring over the next hours and days. Depending on the strength of a given aftershock, its duration and the condition of the damage zone at that time, there can be additional damage to buildings, roadways, gas lines and other structures, as well as additional casualties. After the initial shock and each aftershock impulse, an impulse response follows, an FRP associated with that shock. Thus the time history of the complete set of FRPs associated with an earthquake primary shock and aftershocks can extend over several days, necessitating the modeling of additional phases of earthquake response. Fortunately, aftershocks often account for only a small increment to casualties and damage. This was the case with the Northridge earthquake we are using. Thus we chose to omit any aftershocks in the synthetic earthquake scenario we employ, effectively limiting the simulation to 24 hours and a single FRP.

## 1.2 Domain knowledge

In order to realistically simulate an earthquake event and nominate a particular data fusion system to operate within it, two categories of domain knowledge are necessary. First, the data, models and consequences corresponding to the ground truth of the earthquake itself are needed. How many casualties are created, of what severities and with what geographic distribution? How many hospitals are damaged and to what degree, how many bridges and gas lines? Second, the assets and operational procedures of, and relationships between, the emergency response organizations must be known. Without the first there is no disaster state to be understood, and without the second there is no awareness possible of what the responders need to understand, no knowledge of what constitutes situation awareness for those users.

Domain knowledge concerning the human and structural damage created by the earthquake is derived from HAZUS, a GIS-based natural hazards estimation tool developed by FEMA. By specifying a geographical region, epicenter and severity corresponding to the Northridge earthquake, we can produce a disaster state resembling that event. The input and lay down distributional parameters can be varied to produce a range of earthquake states based on that real world event.

The phenomena of interest to a data fusion system, which processes need to be understood, what constitutes knowledge, must be defined in terms of the user goals and constraints. The end-user of a data fusion capability in this crisis management setting is the first response command network. Domain knowledge of the first response system, how it is structured and how it works, is essential to the fusion design process and realistic modeling of the synthetic environment in which it runs.

Surprisingly, the history of organized multi-agency multi-jurisdictional response to natural disasters is rather short. Following a series of devastating wildfires in 1970 Firescope was formed in California to address the lack of coordination, information sharing or communication standards in wildfire-fighting. There followed the Incident Control System ICS and Multi-Agency Control System MACS, further extending common standards for agencies across the state for a range of natural disasters [6]. In 1994 the state created the Standard Emergency Management System SEMS, organizing all California's emergency offices into a single hierarchical, modular response team [7]. The organizational structure, operational procedures and goals of SEMS inform our synthetic environment and form the framework for defining higher level data fusion hypotheses.

## 1.3 Goals

Within this problem environment a set of goals linked to the Crisis Management Challenge Problem can be explicitly formulated and pursued. These are:

- Define a user-centric distributed high level fusion network for the Northridge earthquake test bed.
- Derive user needs from a work process model within a unified ontological structure.

- Create a needs-driven consistent operating picture (COP) across jurisdictional boundaries.
- Detect and recognize emergent secondary incidents such as HAZMAT spills.
- Provide human decision support through distributed visualization of appropriate elements of the COP to decision-makers in the hierarchy.
- Offload low level decisions to automated processes, but with human supervision and override.
- Define robust metrics for performance and effectiveness, and employ these metrics in the synthetic environment

## 1.4 Higher level fusion

It is widely agreed that the great majority of successful data fusion applications to date have focused on low level (Level 0 and Level 1) capabilities. While effective fusion at the sub-object and object levels offers real performance gains in many of those applications, the creation of dynamic situation assessments and impact predictions associated with higher level fusion are essential fusion products in other potential applications. To encourage a focus on higher level fusion, the AFOSR Challenge Problems specifically call out implementation of systems with Level 2 and higher data fusion capabilities [1].

Higher level fusion is in fact a natural requirement of any effective data fusion system for the Crisis Management Problem. Identification, recognition and attribution of individual objects are not sufficient to guide an effective coordinated disaster response. There is a need to convert data about individual bridges, ambulances, fires, and casualties into usable knowledge about the disaster scene. More abstract entities formed by combinations of aggregated objects and their behaviors must be located in space and time, and attributed. Higher level fusion ontological categories such as process, event and cluster supply essential content for populating a situation assessment which can guide human decision makers and be projected forward in time as impact assessment. Thus we refer to high level fusion in our first stated goal above. Our approach to higher level fusion process and node design will be taken up in Sections 3 and 4 of this paper.

## 1.5 Relation to Robo-Cup Rescue

The Robo-Cup Rescue Challenge, announced by Kitano et al in 1999 [8], was inspired by an earthquake similar to the Northridge earthquake (Kobe City, Japan, 1995, magnitude 7.2). As is the present effort, it is oriented towards disaster mitigation in the FRP. While there are some similarities, there are important differences in scope and goals between this problem and ours.

The Robo-Cup Rescue Challenge incorporates four projects. Three involve building, integrating and deploying real-world systems incorporating robotics and PDA's for disaster rescue. These projects are unrelated to the present effort. The fourth project stipulates development of a comprehensive simulation which can

evolve into a deployable rescue support system in the real world. This is more ambitious than our goals in disaster simulation. Ours is the creation of a plausible synthetic task environment in which data fusion can be integrated and evaluated, finally abstracted and reused. We have no goal of deploying the simulation in the real world. Nor does the Robo-Cup Rescue Challenge have the goal of advancing high level distributed data fusion research.

## 2 The synthetic task environment

The Northridge earthquake of January 17 1994 struck the San Fernando Valley at 4:30AM. Classed as a moderate earthquake of magnitude 6.7, this event caused severe casualty and property damage due to the high population density of this area, which is located within the Los Angeles CA city limits. 72 died and over 1,000 were admitted to hospital, with an additional 9,000 treated and released [9]. Over 12,000 structures were severely damaged, 11 major roadways were closed due to bridge collapses and other structural failures [10].

The ground truth for our synthetic task environment is derived from HAZUS using initialization data that reference this Northridge earthquake. Attributes reported as probability distributions by HAZUS are made definite by performing the indicated probability experiments in order to establish a deterministic ground truth.

The overall software architecture chosen for implementation of this environment is the High Level Architecture (HLA) developed by the Defense Modeling and Simulation Office. HLA is a widely adopted standard for distributed heterogeneous simulation in both the military and civilian communities, and its choice is intended to facilitate reusability.

An HLA-compliant simulation system consists of federates, or separate code modules, interacting via a Runtime Infrastructure (RTI) functional interface. The federates in our synthetic task environment include a ground truth generator, report generator, fusion federate, hospital, dispatch/routing, walk-in and visualization federate. Each models actions critical to simulating a mode of activity in the FRP relevant to casualty outcomes. Report generator collects observations and creates reports which are sent to the fusion federate, which implements all data fusion algorithms and publishes the results to subscribing federates. The hospital federate models the dynamics of hospital medical services, walk-in the dynamics of casualties who choose to seek hospital service without waiting for an ambulance, and dispatch/routing models the assignment of ambulances to casualties and hospitals together with route selection. Visualization delivers visual representations to a human observer or decision maker. Thus for instance ground truth may lay down two casualties at Main and Maple, both of high severity. An observer in the area may incorrectly perceive one severely injured and one lightly injured at that location. A report is sent to fusion, which associates that report with others and perhaps judges there to be two severely injured there. Dispatch/routing then determines how to service these casualties, dispatching an ambulance

and directing it to use a specific route (with alternatives) to the casualty and then on to a specific hospital.

## 2.1 Objects in the environment

Rescue and medical management of existing casualties, together with efforts to mitigate risk of additional casualties, are the dominant goals of early stage disaster response. The objects we model in the synthetic environment are those necessary for these activities to unfold, and the modeled attributes of those objects are those instantiating capabilities linked to the casualty-reduction goals and determining their effectiveness. This does not constitute a comprehensive earthquake simulation in the sense of [4]. Only those objects and attributes most relevant to the FRP goals, and thus most useful in gauging data fusion effectiveness, are represented in the synthetic environment.

Human objects include casualties, police, emergency medical personnel (EMPs) and HAZMAT teams. Casualties are attributed by their ID number, physical description, severity and location. Severities are determined from HAZUS data, as are locations (HAZUS reports location at the granularity of census tract, casualties are then randomly placed within the tract as part of our ground truth lay down). Police serve as observers, cruisers moving from initial lay down locations according to a SEMS-based predetermined damage survey plan. EMPs drive ambulances and deliver medical services as they pick up and transport casualties to hospitals. They also serve as observers. HAZMAT teams respond to hazardous chemical spills caused by rupture of a HAZMAT transporter vehicle in a roadway accident secondary to the earthquake, or rupture of a HAZMAT storage vessel in the damage zone.

Structural objects include hospitals, roadways, bridges and tunnels. Each is attributed by ID number and location. Damage level is associated with hospitals, bridges and tunnels, and link travel times with roadways. Hospital damage, for instance, degrades capacity of that facility. Other structural objects, such as commercial and residential buildings, are not included in the environment. Their damage effects are seen indirectly through the distribution of casualties, which is sufficient for the FRP.

HAZMAT objects include ruptured HAZMAT roadway transporters and ruptured stationary HAZMAT storage tanks. They are attributed by ID number, location, type of hazardous material, and spread of that material.

## 2.2 Reports and Level 1 fusion

Immediately after the primary shock, ground truth is laid down for all objects in the environment. Casualties are characterized and situated, structures tagged by damage level, ambulance and police cruiser initial numbers and locations set.

Initially none of this information is available to the responders. Over logical time reports begin arriving at regional Emergency Operations Centers (EOCs) from observers in the environment: police, EMPs and civilians. Report types include casualty reports and structural damage reports. For instance, an ambulance driver might

report a group of casualties at a certain location as the ambulance heads towards hospital with a full load. A civilian may report that a bridge appears to be severely damaged. Note that civilians are not objects in the environment *per se*. They simply serve as the implied sources of certain reports. Locations and times are selected randomly and casualties or structural damage nearby reported by these civilian reporters.

As in the real world, in this environment reports are uncertain and no observers are completely reliable. The report generation process utilizes confused elements of ground truth to model reports. Each element of ground truth being reported on is subject to a confusion matrix before it is entered a report.

The communications links are also assumed unreliable. Associated with each report is probability of reporting failure and probability of reporting delay. These reports are fused to determine the probabilities of the corresponding object-oriented hypotheses. Association is done through ID numbers, physical descriptions and locations. Associated reports are fused using Bayes algorithms. These results are then used by the situation assessment module discussed in the next 2 sections.

## 3 Design of situation assessment process

The purpose of dynamic situation assessment is to develop probable explanations of the situation based on prior knowledge and incoming transient information. A Situation Assessment (SA) is a stored representation of relations between objects obtained through fusion [12]. The result of situation assessment is a coherent composite picture of the current situation along with prediction of the situation (estimated risk in the case of SA for man-made and natural disasters) to be used by decision makers. In the case of multiple decision makers the situation assessment processes have to deliver a consistent situational picture relevant to each decision maker.

Assessment of the post-disaster situation has specific characteristics, which define requirements for SA architecture and processes. Among these characteristics are:

- Noisy and uncertain dynamic environment with insufficient *a priori* statistical information
- Geographically distributed damage
- Geographically distributed uncertain sources of information often of low reliability .
- Large amount of heterogeneous information
- Resource and time constraints
- High cost of error
- Multiple decision makers with multiple goals and information requirements
- Multiple agencies in multiple jurisdictions

These specific domain characteristics call for a multi-agent distributed dynamic SA process, which has to be adaptive to resource and time constraints, new and uncertain environments and reactive to uncertain inputs. This process also has to accommodate heterogeneous information (both symbolic and numeric).

The SA process exploits reports on casualties and damage of essential facilities, databases, maps, information on prior similar situations, preliminary risk assessment based on historical data and event modeling, and results of domain-specific simulations and models (hospital model, walk-in model, etc.) for creating a dynamic situation picture. The produced situation picture provides the critical characteristics of the state in relation to particular goals, capabilities and policies of the decision makers to serve their ultimate goals, which are to serve the maximum number of casualties, save the maximum number of lives, and reduce risk of additional casualties.

There are three essential components of SA process design. The first component is the Cognitive Work Analysis (CWA) [13], which is a systems-based approach to the analysis, design and evaluation of systems allowing a description of the set of relationships between generic decision tasks, generic activities and available resources. CWA methodology is designed for evaluation of the decision makers' needs to provide understanding of what content various decision makers require from a situation picture and what information should be represented and formulate possible hypotheses about relevant states of the environment.

The second essential component of the SA process is ontological analysis of the specific problem, which denominates the elements of the SA process in terms specific for the disaster domain: objects, attributes, inter-relations, and the dynamic transformations among these objects and relations occurring over time [14].

The third component is a formal SA ontology for catastrophic events, which studies what exists, what can be categorized, and whose goal is to capture the most basic structures of relevant objective reality by developing accurate and comprehensive formal systems that transparently model existing places, times, entities, properties, and relations [15]. The formal ontology framework is necessary to provide a formal structure for ontological analysis of specific type of post-disaster situation, and to assure a certain level of reusability of the designed domain-specific ontology in a different application domain.

The combination of CWA and ontological analysis within the framework of a formal SA ontology is intended to provide sufficient information about the goals, hypothesis, types of objects, relations between them, and processes to support domain specific generation of situational hypotheses and high-level reasoning about these situational hypothesis. The choice of any particular reasoning methods is defined by the domain requirements, the amount of information available, and the level and type of uncertainty of this information. Figure 1 shows the process of SA design.

Situation in the first phase post-earthquake scenario consists of a set of elementary situations and their compositions. Elementary situation nominations are based on the results of the CWA and correspond to essential elements of information required by decision makers for taking actions. Among elementary situations to be considered are Communication system situation, Transportation system situation, HAZMAT situation

(secondary threat), Casualties situation at different levels of aggregation, Hospital situation at different levels of aggregation, Ambulance situation, and Resource situation. Each elementary situation, when considered at a certain point of time (current or future), is an event, represented by a set of hypothesis with confidence levels, by risk associated with this situation, and a set of attributes with their values characterizing this situation. Over time, each situation is also a process, which is characterized by behavior of its attributes.

## 4 Situation assessment node

The SA design process architecture is presented in Fig 1. Information is evaluated to produce a consistent decision state estimate, which is presented for the system application (either the next automated steps or for presentation to a user).

### 4.1 Preprocessing

Situation assessment further processes and aggregates information about objects obtained as the result of level 1 fusion of reports on casualties and facilities damage. The results of SA depend heavily on the quality of results of the Level 1 fusion processes. Although report fusion runs constantly, fused information on each particular casualty or structure cannot enter the SA process until the quality of this information is sufficient to ensure the quality of the resulting SA. At the same time the SA process cannot wait until the stream of reports about a certain casualty or structure is complete. Waiting may result in unacceptable decision latency, leading to either wasted resources or lost lives.

This situation calls for preprocessing, which can be implemented as decision making under a time constraint. Report fusion on a certain object should be stopped and the results passed to the SA process either when the quality of the level 1 process is acceptable or a certain deadline has been reached. The quality of the level 1 estimation can be assessed, for example, by comparing the confidence level of estimates with a time-varying threshold.

### 4.2 Situational state estimation

Formally, let  $\Omega$  be a set of possible states of the environment,  $\Omega^k \subset \Omega$  a set of possible states of the environment relevant to decision maker  $k$ , and  $\{S^k(t)\} \subseteq \Omega^k$  is a situational picture relevant to decision maker  $k$  at time  $t$ . Set  $\{S^k(t)\}$  is represented by a set of pairs  $\{H_i^k(t), Bel_i^k(t)\}$ , where  $\{H_i^k(t)\}$  are hypotheses of decision maker  $k$  at time  $t$  and  $\{Bel_i^k(t)\}$  are corresponding levels of confidence into each hypothesis.

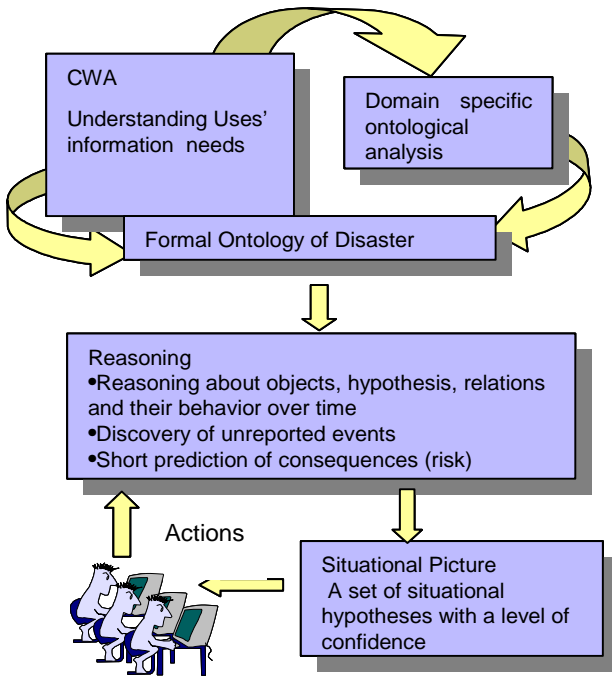


Figure 1. Situation assessment design process

### 4.3 Modeling framework

The modeling framework selected for our system represents a combination of domain specific models such as a hospital model and a dynamic dispatch/routing model, with time-dependent belief networks. Time-dependent belief networks (BN) are graphical models representing causal and belief relations among random variables and give an option to update those beliefs upon arriving new information. Belief networks provide intuitive and causal representations of real-world applications, and are supported by a rigorous theoretical foundation [16]. They allow expert knowledge and empirical observations to be combined (fused reports in our case), and to provide efficient uncertainty representation, which make them applicable for SA. BNs consist of two parts: a directed acyclic graph representing qualitative relations between random variables, and a set of a priori and conditional beliefs which quantify these dependencies. Building the graphical representation and modeling a priori and conditional beliefs present the major challenges of BN. In our system, the graphical representation is derived from the SA ontology [17].

A priori and conditional beliefs in BNs in most cases are expressed in the framework of probability theory (as Bayesian Networks) and learned from historical data and from expert knowledge and databases. In the natural and man-made disaster context, numeric historical data is sparse, the uncertainty, vagueness, and imprecision of attributes, properties, and relations are high, and many relations are represented in vague symbolic linguistic ways (*high, close, soon*, etc). All this makes the task of deriving useful model probabilities very difficult. One of attractive ways of dealing with this highly uncertain and incomplete environment is to consider here a combination of a qualitative belief network with a Bayesian network

[18]. This combination provide probabilistic reasoning in a qualitative way when numerical probabilities are not available.

### 4.4 Belief change

An important component of SA is checking consistency of the situation picture in a Belief Change process. Situational state estimation at time  $t$  consists of updating the estimates obtained at time  $t-1$ . It starts at time  $t=1$  with updating and revising the a priori situation estimation, which is based on domain knowledge about initial risk. For example, in the case of the earthquake, the initial SA is based on modeling, seismic and geophysical information about the severity of the earthquake, knowledge about building structure and vulnerabilities, population densities, and other specific and relevant information concerning the disaster zone.

New information obtained at time  $t$  drives changes to the state estimates obtained at time  $t-1$ . Traditionally, the nature of information combination in such cases is considered non-symmetrical and new information is given priority to existing information while accounting for reliability of this new information (see, e.g. [19-21]). However, in the distributed case, we need to consider separately two process, *belief revision* and *belief update*, which treat priority of incoming information differently [22].

The belief revision process modifies existing estimates at  $t-1$  based on new information obtained at time  $t$  to refine the SA at time  $t-1$ . i.e., belief revision refers to a static situation, although it can be used in a dynamic situation when referred to locally stable conditions. The belief update process, on the other hand, modifies existing estimates at  $t-1$  based on new information obtained at time  $t$  to build a new SA at time  $t$ .

Belief revision decides what beliefs (old or new) should be discarded to accommodate new information. Revision in the static case is based on conditioning while reliability of all beliefs has to be taken into account (see, e.g. [23]). New information may be discarded if it contradicts either domain knowledge or totally reliable previous information.

In dynamic situations, incoming information describes the changed situation and the nature of belief combination is not symmetrical. In such situations *belief update* has to be considered. In belief update an agent's beliefs should be adjusted to be consistent with a priori knowledge as well as knowledge concerning new events which occurred in the changing problem environment. Belief update attempts to decide what changes in the world led to this new information. Here incoming information is given higher priority provided that its reliability is taken into account. Transition from  $Bel(t-1)$  to  $Bel(t)$  should obey the principle of minimal change of previous beliefs to make it compatible with the new information. In dynamic situations, a Kalman-like approach to belief update can be adopted ("model-based" BR) [24]. In this case revision consists of a prediction step based on a selected model of the evolution of the world and a revision step, in which predicted state of the world is modified based on incoming

information while taking into account its reliability. Incoming information can be rejected if this new state deviates too far from the predicted state. In our system the consistency check will be based on a priori risk estimation, relation between fused reports of different types and database information, and may be different for different elementary situations.

## 4.5 Decision state estimation

First response phase casualty mitigation operations are under severe time and resource constraints, and timely decision making and swift action are required. At the same time the cost of false alarms can be very high since valuable resources might be diverted from the location where it later becomes clear that they are critically needed. The cost of waiting for additional information, or cost of additional computation delay, has to be justified by the benefits of achieving a more accurate SA. Therefore, as in the preprocessing step described in Section 4.1, the result of aggregation and SA should be determined to be

of a minimum threshold quality before being allowed to be used by other processes or passed on to decision makers. The state of acceptable quality is known as a "decision state". The process of decision state estimation requires criteria for defining situation quality. One of the ways of dealing with this problem is to select a set of pivotal situational hypothesis and then to define a quality of the situation containing this hypothesis by a time-dependent confidence level associated with this hypothesis. The nomination of the pivotal hypotheses and a time-dependent confidence level can be obtained as the result of CWA. In certain situations, when decisions based on the resulting decision state estimations have very serious consequences, a sensor management process can be employed. For example, a highly reliable sensor, perhaps a policeman or structural engineer, can be tasked to observe the situation in question. The situation assessment processing logic is illustrated in Figure 2 which is located after Section 6 References.

## 5 Conclusions

A synthetic task environment in the context of natural and man-made disasters has been constructed to explore data fusion system design and performance. Specifically, the first response phase of an event similar to the Northridge CA 1994 earthquake is modeled. Secondary HAZMAT spills caused by the primary earthquake shock are included. Key features of the environment include incorporation of higher-level and distributed fusion capabilities, and surveillance for secondary incidents, which may need to be inferred rather than directly reported to the fusion nodes.

A methodology for designing an SA process and architecture of an SA node has also been introduced. The methodology is based on the combination of cognitive work analysis and ontological analysis within a framework of a formal situation assessment ontology. The proposed SA process architecture represents a multi-step situation estimation, in which the real-time initial situation is estimated and evaluated to produce estimation of a consistent decision state estimate. The evaluation comprises a quality step, in which both output of the level 1 fusion process and the current estimate are evaluated to see if they satisfy certain quality criteria, and a consistency evaluation step, in which the decision state estimate is examined for consistency in a belief change process. The proposed architecture is not application specific and can be a component of a general fusion processes [26].

The next phase of this effort will be directed towards continued algorithm development and the systematic testing of the system, addressing the question of nomination and analysis of performance metrics. Particular attention will be paid to the specification of measures of effectiveness which guide optimization of the higher-level fusion performance.

## Acknowledgements

This work was supported by the AFOSR under award F49620-01-1-0371. The authors gratefully acknowledge valuable input from Dr. John Tangney, Dr. John Salerno, Dr. James Llinas, Dr. Ann Bisantz, Dr. Chris Bowman and Dr. Eric Little.

## 6 References

- [1] John F. Tangney. Programs in Higher Levels of Information Fusion. *IEEE Aerospace and Electronics Systems Society Magazine*, pages 21-25, Nov. 2003.
- [2] James Llinas. Information Fusion for Natural and Man-Made Disasters. In *Proc. Fifth Int. Conf. Information Fusion*, pages 570-574, Annapolis, MD, USA, 8 July-11 July 2002. Int. Soc. Information Fusion, Sunnyvale, CA 2002.
- [3] U.S. Air Force Scientific Advisory Board Committee on Building the Joint Battlespace Infosphere. Building the Joint Battlespace Infosphere. Report *SAB-TR-99-02*. Washington D.C., 7 December 2000. U.S. Air Force Scientific Advisory Board, 2002.
- [4] Ikuo Takeuchi, Shigeru Kakumoto and Yozo Goto. Towards and Integrated Earthquake Disaster Simulation System. *Proc. First Int. Workshop on Synthetic Simulation and Robotics to Mitigate Earthquake Disasters*. Padova, IT, 5 July-8 July 2002, Universita di Roma 2002.
- [5] John Salerno and Paul Bello. Information fusion. Information Institute General Workshop II. Rome NY 10 June-11 June 2003.
- [6] History of ICS. *Incident Command System National Training Curriculum Report*. Washington DC, October 1994. National Wildfire Coordinating Group, 1994.
- [7] Gray Davis and Dallas Jones. *SEMS guidance for special districts*. July 1999, Fresno CA. Standardized Emergency System Advisory Board.
- [8] Hiroaki Kitano et. al. *Robo-Cup Rescue: search and rescue in large scale disasters as a domain for*

- autonomous agents research*. IEEE Int. Conf. on Sys. Man and Cyb., pages 739-743, Tokyo, October 1999.
- [9] Michael E. Durkin. Fatalities, nonfatal injuries, and medical aspects of the Northridge earthquake. Mary C. Woods and W. Ray Sieple (eds), The Northridge CA earthquake of 17 Jan 1994. California Dept. of Conservation Div. of Mines and Geol. Spec. Pub. 116. pages 247-254, 1994.
- [10] Earl Aurelius, ed. *The January 17 1994 Northridge CA earthquake*. Houston TX 1994. Tech rept of ABS Consulting Inc., 1994.
- [11] Qiang Gong, Arun Jotshi and Rajan Batta. Dispatching/routing of emergency vehicles in a disaster environment using data fusion concepts. *submitted to the Seventh Int. Conf. Information Fusion* Stockholm SW 28 June–1 July 2004.
- [12] D. Lambert. Grand Challenges of Information Fusion. *Proc Sixth Int. Conf. Information Fusion*, pages 570-574, Cairns, Australia, pages 213-220, 8 July–10 July 2003. Int. Soc. Information Fusion.
- [13] A. Rasmussen, J. Pejtersen, L. Goodstein. *Cognitive Systems Engineering*. Wiley, New York, 1994.
- [14] Wayne Johnson, Ian D. Hall. From Kinematics to Symbolics for Situation and Threat Assessment. *Proc. Conf. on Information, Decision and Control*. Adelaide, Australia, 8 Feb.-10 Feb. 1999. Defence Science and Technology Organization, 1999.
- [15] Eric Little. Foundations of Threat Ontology (ThrO) for Data Fusion Applications. *Center for Multisource Information Fusion Technical Report*, Buffalo NY, 2003.
- [16] Finn. V. Jensen. *An introduction to Bayesian Networks*, Springer, New York, 1996.
- [17] Eveline. M. Helsper, Linda. C. van der Gaag. Building Bayesian networks through ontologies. F. van Harmelen (ed.). *Proc. 15th Euro. Conf. on Art. Int.* pages 680-684, Amsterdam, the Netherlands, IOS Press, 2002.
- [18] Thuong Doan, Peter Haddawy, TienNguyen, and Deva Seetharam. A Hybrid Bayesian Network Modeling Environment. In *The Nat. Comp. Sci. and Eng. Conf. NCSEC* 1999.
- [19] Peter Gärdenfors. *Belief Revision*. Cambridge University Press, Cambridge, U.K., 1992.
- [20] Craig. Boutilier, Nir. Friedman, Joeseeph. Halpern.. Belief Revision With Unreliable Observations. *Proc. of the Fifteenth Nat. Conf. on Art. Int.* pages 127--134, Madison, WI, 26 July-30 July, 1998. American Assoc. for Art. Int. Menlo Park CA USA, 1998.
- [21] Didier Dubois and Henri Prade. Introduction: Revision, Updating, And Combining Knowledge. *Handbook of defeasible reasoning and uncertainty management systems Vol. 3*. Kluwer Academic Publishers, London, U.K., 1998.
- [22] Isabelle Bloch, Anthony Hunter, Introduction, Fusion: General Concepts and Characteristics, *International Journal Of Intelligent Systems*, Vol. 16, 1107–1134, 2001 John Wiley & Sons, Inc, 2001
- [23] Aldo Dragoni. Belief Revision: From Theory To Practice. *The Knowledge Engineering Review*. Vol. 12(2), Cambridge University Press, 1997.
- [24] Salem Benferhat, S., Didier Dubios and Henri Prade. Kalman-like Filtering in a Possibilistic Setting. in the *Proc. of the Euro. Conf. on Art. Int.* 2000. Berlin GE, 20 Aug.-25 Aug. 2000. European Coordinating Committee for Art. Int., West Lothian, U.K., 2000.
- [25] James Llinas, Christopher Bowman, Galina Rogova, Alan Steinberg, Edward Waltz, and Frank White, Revisions to the JDL Data Fusion Model II, in the *Proc. the. Seventh Int. Conf. Information Fusion* Stockholm SW 28 June–1 July 2004. Int. Soc. Information Fusion.

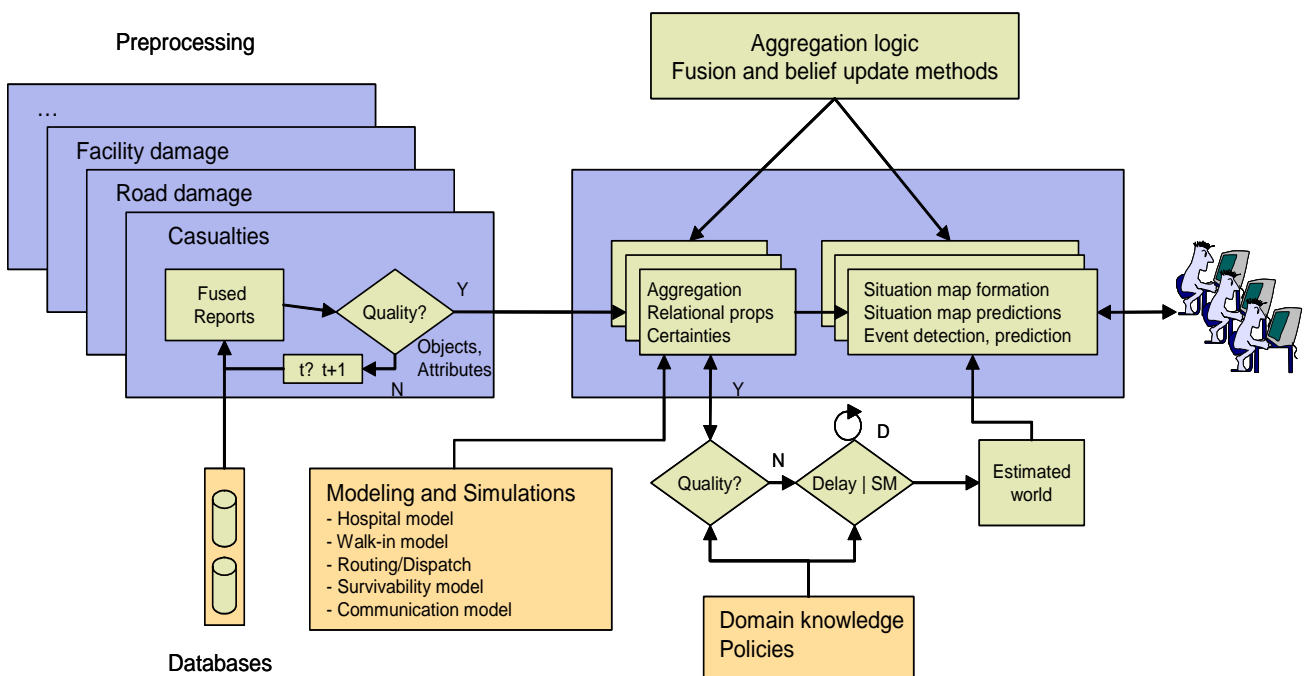


Figure 2. Situation assessment processing