Distributed Perception Networks for Crisis Management

Marinus Maris
University of Amsterdam
mmaris@science.uva.nl

Gregor Pavlin
University of Amsterdam
gpavlin@science.uva.nl

ABSTRACT

Situation assessment in crisis management applications can be supported by automated information fusion systems, such as Distributed Perception Networks. DPNs are self-organizing fusion systems that can infer hidden events through interpretation of huge amounts of heterogeneous and noisy observations. DPNs are a logical layer on top of existing communication, sensing, processing and data storage infrastructure. They can reliably and efficiently process information of various quality obtained from humans and sensors through the existing communication systems, such as mobile phone networks or internet. In addition, modularity of DPNs supports efficient design and maintenance of very complex fusion systems. In this paper, a fully functional prototype of a DPN system is presented that fuses information from gas sensors and human observations. The task of the system is to compute probability values for the hypothesis that a particular gas is present in the environment. It is discussed how such a system could be used for crisis management.

Keywords

Crisis management, multi-agent systems, distributed perception, information fusion

INTRODUCTION

In crisis situations, decision makers must respond appropriately to different adverse situations. Effective mitigation of hazardous situations, on the other hand, requires profound knowledge of the relevant events in the affected environment. However, often crucial events cannot be observed directly. Consequently, such "hidden" events must be inferred through the appropriate interpretation of information sources, such as sensors and humans. We call such interpretation information fusion.

In general, information fusion processes range from the interpretation of sensor data to the interpretation of complex features of in order to present useful information to human operators. Such gradual interpretation is captured by two most influential models, namely the Joint Director's of Laboratories (JDL) data fusion model (Llinas, 2002) and the Endsley model (Endsley, 2005). Both the JDL and the Endsley model, emphasize the ability to fuse relevant features and present the operator with the relevant information only.

In this paper, the focus is on a part of such complex systems, namely perception. Typical approaches to perception for crisis management take a rather centralized view, in which all information that is available is sent to the operators. However, these centralized approaches suffer from several problems. They cannot cope with the growing complexity of problems and they cannot fully exploit the potential of modern communication and sensory technology. Namely, in modern applications huge amounts of information are sent to a central office which can result in clogging of communication resources, while operators must make sense out of huge amounts of information.

Moreover, due to the recent advances in communication and sensing technology we can obtain large amounts of relevant information which simply cannot be processed manually within the required time constraints. In particular, the existing mobile phone networks can provide huge amounts of information that can be very valuable in a crisis situation and should be exploited to the maximum (IFRC, 2005). This requires an appropriate information processing system which can extract the relevant bits of information. In addition, in crisis management applications situation assessment must be quick and reliable, since delayed or misleading assessment can lead to inappropriate action selection which, in turn, may have devastating consequences on the further course of events.

In order to be able to deal with these challenges, we take the stance that situation awareness in a crisis situation can be improved by technological support systems (Burghardt, 2004). With respect to this we propose Distributed Perception Networks (DPN) as a tool to handle the vast number of information flows stemming from sensors and humans as occurring during a crisis. In particular humans can be viewed as useful and versatile sources of information because they are good interpreters and can recognize a rich spectrum of phenomena. The DPN approach is relevant for a significant class of crisis management applications, since it supports:

- Quick delivery of the relevant information;
- Discovery and integration of all relevant information sources at runtime;
- Inference based on uncertain models and noisy information;
- Simple design and maintenance of complex fusion systems;
- Efficient processing and usage of the existing information sources.

We focus on perception which is based on information processing at different abstraction levels of complex fusion systems. The fusion system automatically contacts sensors and humans and incorporates their observations into the fusion results. DPN systems can send SMS messages to humans in the affected area to learn about their observations. Another important contribution of our method is the completely distributed approach of the fusion processes. In particular within the scope of crisis management, distributed information is crucial to avoid a centralized system, which introduces a single-point-of-failure and is vulnerable to sabotage (Turoff et al., 2004). By using DPNs, information processing is not only performed at the control room, but mostly near the information sources themselves. In this way the DPNs support quick response times and graceful degradation (Pavlin et al., 2004).

DISTRIBUTED PERCEPTION NETWORKS

As mentioned before, large amounts of heterogeneous information as well as many required inference steps make centralized approaches to information fusion impractical. We approach these problems with Distributed Perception Networks (Pavlin et al. 2004). A DPN is a multi-agent system (Wooldridge, 2002) in which each agent takes care of partial fusion. We distinguish sensor, human and fusion agents. Sensor agents are created by "wrapping" a sensor with agent-specific capabilities, such as local information processing, communication and negotiation. Human agents take care of bi-directional communication with humans via the cell-phone, in this case through SMS. The fusion agents take the input from sensor and human agents for further information fusion. Such a group of agents constitute a particular DPN. Each DPN is specialized for a specific fusion task, i.e. it estimates the probability distribution over a single variable, e.g. the presence of gas X. Thus, information fusion in DPNs is a hierarchical process where different participating agents cooperate in order to gradually map observational evidence to the probability distributions over the hypotheses of interest. Also, an agent-based approach to fusion provides the means to keep partial processing close to the spatially dispersed information sources, which reduces the danger of communication and processing bottlenecks.

We stipulate further that the fusion structure of a DPN must be based on a causal model, since causal processes can capture the knowledge of domain experts in a very intuitive manner (Pearl, 1999). A popular method to describe causal processes is through Bayesian Networks (BNs). With BNs we can express causal relationships and formulate probability values over an arbitrary combination of states of variables in an efficient way (Jensen, 2001). Moreover, BNs support systematic modeling of causal processes involving rather heterogeneous concepts. In other words, BNs allow processing of very heterogeneous types of information in a uniform and efficient way.

A typical example of a DPN is shown in figure 1. This DPN evaluates the hypothesis *Ammonia*, i.e. it calculates probability values for its top node. Each dotted rounded rectangle represents one DPN agent, running a local BN as part of the entire BN to evaluate the hypothesis. There are agents processing sensor measurements (*Scon1* and *Scon2*), agents accessing humans (*CyanMD, CyanCV, Smell, NauseaMD,* and *NauseaCV*), and fusion agents (*Cond, Cyan, Smell* or *Nausea*). It is assumed that the agents are distributed throughout a network of computing devices. The top node *Ammonia* communicates with the crisis center to report the results of the probability calculations.

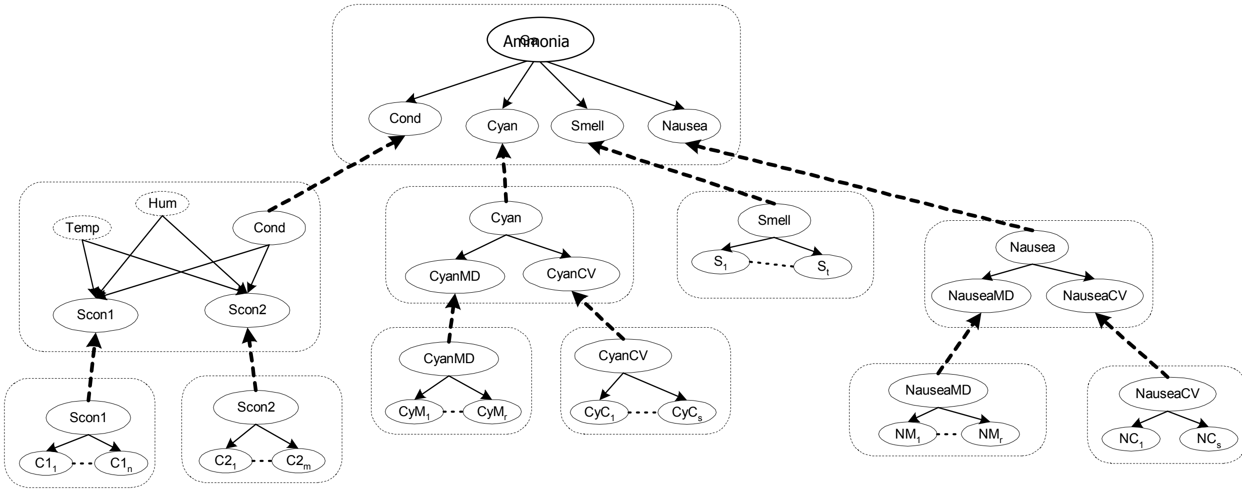

Figure 1. A DPN fusion organization where each dotted rectangle represents a DPN agent. Thick dashed lines represent communication between cooperating agents, which share partial fusion results. Each agent makes use of a local Bayesian Network, a probabilistic model that captures specific expertise over a certain domain. The agents assemble their local models into a distributed causal model that supports evaluation of the hypothesis about the existence of "Ammonia".

We developed a DPN prototype environment and created a system which implements the DPN shown in figure 1, including gas sensor data processing and automated bi-directional SMS messaging. The theoretical aspects of DPN are extensively discussed in (Pavlin et al., 2004; Nunnink and Pavlin, 2004/2005; Oude et al, 2004). For the rest of this paper, the emphasis will be on a deployment example of DPN in a crisis management environment.

USE OF DPN IN CRISIS MANAGEMENT APPLICATIONS

In this section we use a simple example to explain how a DPN system can support quick and reliable situation assessment in a crisis management application by harnessing the potential of the existing sensory and information infrastructure.

Let's assume that ammonia escaped from a ship in a busy harbor and a toxic cloud is forming over a densely populated neighborhood. Clearly, the decision makers must learn about the fact that ammonia has escaped as quickly as possible in order to be able to scale up the relevant organizations and plan appropriate measures. However, the presence of ammonia cannot be observed directly. Instead, the existence of a high concentration is estimated through the interpretation of different observations about the symptoms that ammonia may cause, such as sensor reports, typical smell as well as health problems.

Some sensing capabilities might be provided through a network of ammonia sensors dispersed throughout the harbor. We further assume that a gas sensor located in the ship's vicinity produces an initial report indicating an unusual concentration of ammonia (the "Gas Sensor" in figure 2). Typically, such sensors are relatively unreliable and can often produce either false alarms or might not be able to detect a critical state. On the other hand, the decisions based on such observations are mission critical as a false alarm may negatively influence the further course of crisis response.

In other words, perception assessments must be reliable. Therefore, we want to support the assessment robustness by considering relevant information obtained from other types of information sources. For example, the decision makers could contact fire fighters, paramedics and police squads that routinely keep watch of the harbor and ask them about the symptoms, such as smell, health problems, etc. However, a manual process of querying people in the field is not trivial and can be very inefficient as contacting the relevant persons takes time and the interpretation of different symptoms requires significant domain expertise.

We can cope with these problems by using a system of DPN agents dedicated to the reasoning about the existence of toxic gases in the area of interest. Each sensor is wrapped into a DPN agent and there exist DPN fusion agents, each having a limited domain expertise in toxic gases and the symptoms they cause. DPN agents automatically form organizations, each geared towards determining a certain type of a gas. Such a system of agents provides a "bridge" between many low-quality information sources and the crisis managers in the control room (see figure 2), where the DPN-agents are indicated by the gray squares, each running a local BN (depicted by the orange ellipses). Such a system distills relevant information (e.g. the probability that *Ammonia* is present) from large amounts of "low level" information.

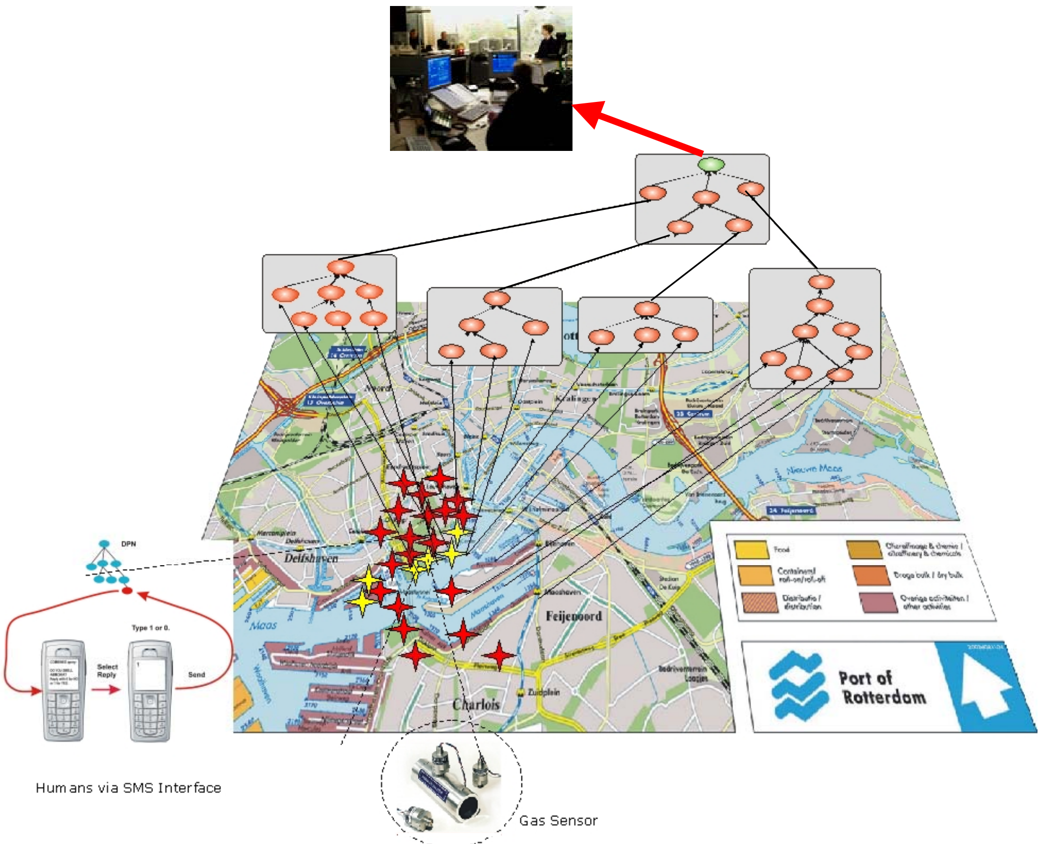

Figure 2. Crisis management support system with DPNs

In a system of DPN agents, an observation by an unreliable ammonia sensor initiates a self organization process resulting in a specific DPN, i.e., an organization of agents specialized for reasoning about the existence of ammonia (see figure 1). During the self-organization process, DPN agents quickly discover all available information sources that could provide information relevant for the estimation of the presence of ammonia. This is possible because DPN agents support reasoning based on rigorous causal models. Such causal models explicitly encode causal relationships between phenomena (e.g. existence of high ammonia concentration) and their typical symptoms. Consequently, agents can exploit such causal models for a systematic discovery of the relevant information sources and their integration into meaningful information fusion systems. For example, given a high ammonia concentration, also reports from other types of ammonia sensors in the ship's vicinity should be considered. In addition, the causal model could encode the fact that when people are exposed to ammonia, they can smell it and might develop typical health related symptoms, such as headache, cyanosis, short breath, etc. In such a case, agents can check the symptoms encoded in the model and, for each symptom, form simple queries that can be broadcast to the people in the affected area. For example, after an ammonia sensor reported the presence of a gas, the DPN agent that has access to the mobile phone network would broadcast the following SMS messages to all policemen, paramedics and firefighters in the area (see figure 3):

"Do you feel Nausea? Reply with 0 for NO or 1 for YES"

If a person replies to such a message, the reply is sent back to the DPN agents which use it as evidence in the BN. Similarly, DPN agents, reasoning about the existence of ammonia, could query medical staff and people in the street about a variety of symptoms. In this manner, we simplify communication between agents and humans. Namely, the DPN agents guide the responding process such that humans provide only information about the facts that are relevant for reasoning, in this example, about ammonia. In addition, by limiting the responses to a simple "yes/no", we avoid the necessity of parsing natural language, thus reducing potential ambiguity in communication. Note further that due to the existing mobile phone infrastructure, we can often quickly access many people in the area and obtain large amounts of relevant information. This allows very robust assessment even when we use low-quality information sources (Nunnink and Pavlin, 2006).

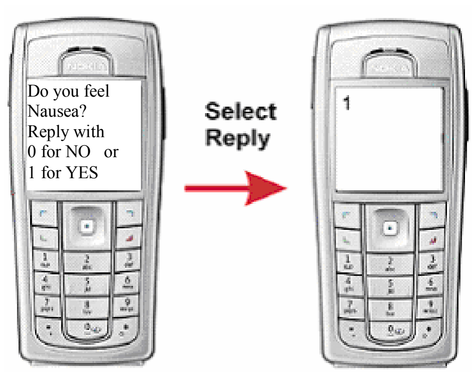

Figure 3. Typical SMS-message for use in the DPN reasoning system

CONCLUSION

This paper presents an approach to information fusion as a method to support situation assessment systems. An important aspect of our system is that it exploits the existing mobile phone infrastructure to contact humans. A prototype system was presented that actually interacts with sensors and humans using local fusion models and multi-agent technology. The seamless integration of information both from sensors and humans suggests that this approach may be useful in a significant class of crisis management applications, such as detection and monitoring of toxic gases, forest fires, tsunamis, adverse meteorological phenomena, etc. Currently, we are looking into possibilities for cooperation with governmental institutions to perform pilot studies in order to test and evaluate the capabilities of this approach in real world applications.

ACKNOWLEDGMENTS

This work was carried out as part of the project Combined Systems within the Decis laboratory in Delft in The Netherlands. We thank our colleagues from the University of Amsterdam for their academic support and in building components for the DPN system.

REFERENCES

1. Burghardt, P. The Combined Systems Point of View. In *Proceedings of ISCRAM*, Brussels, May 2004.
2. Corkill, D. Representation and Contribution-Integration Challenges in Collaborative Situation Assessment. In *Proceedings of the Eighth International Conference on Information Fusion*, Philadelphia, Pennsylvania, July 2005.
3. Endsley, M. R. Toward a theory of situation awareness in dynamic systems. *Human Factors, 37(1)*, 1995. pp. 32-64.
4. IFRC (International Federation of Red Cross and Red Crescent Societies). World Disasters Report 2005.
5. Jensen, F.V. Bayesian Networks and Decision Graphs. Springer. 2001.
6. Llinas, L., Bowman, C., Rogova, G., Steinberg, A., Waltz, E. & White, F., Revisiting the JDL Data Fusion Model II. *Fusion2004,* Stockholm, Sweden, pp1218-1230, 2004.
7. Llinas, J. 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.
8. Nunnink, J. and Pavlin, G. A probabilistic approach to resource allocation in distributed fusion systems. In *Proceedings of AAMAS*. Utrecht, The Netherlands, 2005. pp. 846-852
9. J. Nunnink and G. Pavlin. Towards Robust State Estimation with Bayesian Networks: A New Perspective on Belief Propagation. In *Proceedings of IAS-9,* Tokyo, 2006.
10. Oude, de P., Ottens, B. and Pavlin, G. Information Fusion with Distributed Probabilistic Networks. In *Proceedings of Artificial Intelligence and Applications,* 2004. pp. 195-201.
11. Pavlin, G., Maris, M., and Nunnink, J. An Agent-Based Approach to Distributed Data and Information Fusion. In *Proceedings of Intelligent Agent Technology*. Peking, China, 2004
12. Pearl, J. Causality: Models, Reasoning and Inference. Cambridge, 1999.
13. Scott, P.D. and Rogova, G.L. Crisis Management in a Data Fusion Synthetic Task Environment*. In Proceedings of FUSION 2004*. Stockholm, Sweden, 2004.
14. Turoff, M., Chumer, M., Yao, X. and Van de Walle, B. "The design of a Dynamic Emergency Response Management Information System", *Journal of Information Technology Theory and Applications* (2004).
15. Wooldridge, M. An Introduction to Multiagent Systems, John Wiley and Sons, 256 pages, March 2002. ISBN: 0-471-49691-X.