# Coordination of First Responders Under Communication and Resource Constraints

# (Short Paper)

Robert N. Lass Joseph B. Kopena Evan A. Sultanik Duc N. Nguyen

Christopher P. Dugan Pragnesh J. Modi William C. Regli

College of Engineering
Drexel University
Philadelphia, PA 19104
{urlass, tjkopena, eas28, dn53, cjd48, regli}@cs.drexel.edu

## ABSTRACT

This paper discusses the application of distributed constraint optimization to coordination in disaster management situations under sub-optimal network conditions. It presents an example system for the problem of shelter assignment and outlines some of the challenges and future research directions that must be addressed before real-world deployment of distributed constraint optimization becomes a reality.

## Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*; J.7 [**Computers in Other Systems**]: Command and control

## Keywords

Distributed constraint reasoning, evacuation, MANETs

## 1. INTRODUCTION

Coordination among emergency personnel and organizations is a critical factor in the successful management of any natural or other disaster. Authorities must assign tasks, distribute resources such as food, medical, and shelter space, negotiate frequencies, and so on. Effective coordination ensures efforts are not duplicated and all resources—including time—are used well. Unfortunately, accomplishing this on a large scale in real time is currently difficult at best, generally relying on much manual problem solving and communication over unreliable and limited analog voice radios. Automated, networked systems stand to offer much needed capability to augment human decision making and enable better coordination in managing disaster scenarios.

Our belief is that the majority of this decision making required of first responders in managing an emergency situation may be seen as the propagation and solving of constraint reasoning problems. Critical to this is the sharing of information and derivation of solutions that maximize global utility, rather than limited local decision making. Achieving this under challenging, non-traditional networking conditions requires development and enhancement of coordination algorithms. The act of coordination itself must be network aware and adaptive to dynamic characteristics of the underlying network, as well as built on and tailored for robust communications primitives.

Continually improving the speed and effectiveness of modern, large-scale, multi-organization disaster response require automated coordination between organizations and personnel in these and many other decision areas. Computer supported techniques must be employed to share and manage the flood of information available, make decisions under tight deadlines, and operate under the external stresses imposed in such environments. Without such support, geographically and organizationally dispersed responders will not be able to develop a shared understanding of the situation at hand to make globally optimal decisions.

Wireless networking technologies are an essential foundation of such support. Mobility needs and the absence of infrastructure require systems built on digital, wireless communication. To be practical in disaster settings, situational awareness, command and control, and coordination applications must rely on rapidly deployable and ad-hoc, wireless mobile computing networks built as "systems-of-systems," encompassing multiple types of communications hardware; heterogenous nodes such as PDAs, laptops, and cell phones; and a wide variety of applications. Such networks, however, present many challenges to information sharing and coordination due to battery power, limited processing power, latency, bandwidth, and network connectivity.

The principle contributions of this paper are:

- A description of wireless communication constraints that effect coordination of emergency responders;
- A new application of distributed constraint optimization to addressing coordination under these challenging networking settings; and
- An example of an evacuation coordination scenario and prototype application.

The paper is organized as follows: An example scenario and coordination problem is presented in the following section. Network characteristics in these settings are then outlined,

**Cite as:** Coordination of First Responders Under Communication and Resource Constraints (Short Paper), Lass, *et. al.*, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008,Estoril,Portugal,pp.1409-1412.

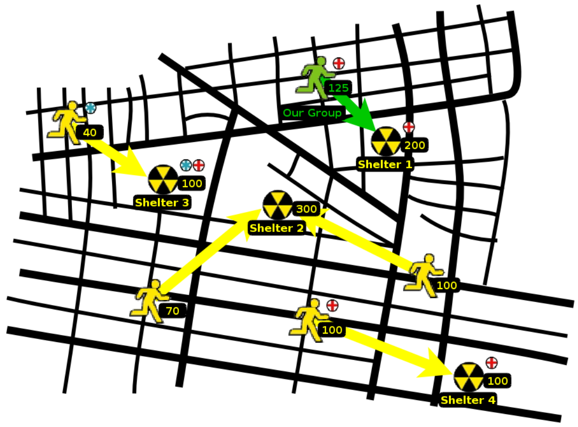


**Figure 1: Example neighborhood sheltering scenario displaying shelter assignments.**

followed by a formalization of distributed constraint optimization and revisiting of the example application. Lastly, we discuss the suitability of distributed constraint optimization for these settings, networking characteristics of several algorithms, and conclusions.

## 2. SCENARIO: DISASTER EVACUATION

Evacuation of neighborhoods, cities, and regions is a major component of responding to many emergency situations, natural or otherwise. The intuitive and easy response of moving evacuees to the closest refuges can quickly overwhelm shelter capacities. Insufficient knowledge and decision processes may also lead to mismatches between evacuees' needs and shelters' capabilities, such as medical facilities. Our work addresses this problem by developing tools and techniques to help emergency personnel create a shared and accurate understanding of the situation, and coordinate to make the best decisions for the group.

A central premise is that there are emergency personnel leading groups of evacuees to available shelters. These authorities are equipped with handheld, networked devices to monitor and coordinate actions. Figure 1 depicts a notional example scenario. There are groups of people represented by the walking figures and available shelters (the fallout shelter symbols) within the local area. Each group has several traits such as size and medical needs. Shelters mirror these with capacity and medical capabilities (the red and blue icons attached to the shelters). The problem is that of assigning groups to shelters in a globally optimal fashion, i.e. not overcrowding any shelter or wasting medical resources.

Our example system for this task is built on handheld computing devices—tablets and PDAs—communicating wirelessly over a mobile, ad-hoc network. Wi-Fi networking is used for testing and demonstration purposes, with true usage most likely requiring different, longer-ranged radios, such as WiMAX. It assists first responders by informing them of the current situation (shelter and group sizes), sharing situational data, aiding decisions making, and helping the users monitor these decisions as events unfold.

The application of planning to evacuation operations has been studied in several projects, such as [8]. Most work in this area is focused on centralized, *a priori* development of plans and procedures. In contrast, the work here focuses on supporting the actual conduct of such operations via situational data exchange and distributed decision making aids. These are largely complementary areas. A priori plans are an assumed input here, *e.g.* candidate routes and destinations, while emergency response planning tools may make use of systems such as that presented here to monitor execution and conduct decision-making.

Work also exists on the use of sensor networks for emergency navigation, such as [13]. The focus is generally on sensing hazards such as fire and gases, determining a route around the obstacle, and using environmental signals such as lights to safely route people toward safety. This is also largely complementary work to that presented here. Such navigation systems are generally based distributed path planning or network routing and operate at a low level of detail, navigating around obstacles. The work here aims to reason on more abstract properties such as available space and medical capabilities, which could potentially then rely on emergency navigation systems to help carry out decisions.

## 3. NETWORKING CHALLENGES

A major challenge in this work is the development of coordination techniques that are robust and adaptive on realistic wireless networks. Infrastructure-free and mobile networks present significant problems to systems and application developers. Software cannot be developed in the same manner as for infrastructure-based networks and be expected to perform and behave similarly. This is due to a number of distinctive characteristics of networking in disaster settings. In [3], eight properties were identified that affect applications on MANETs in four different ways:

- **Latency:** Due to a higher rate of link errors, MANET specific service models and other factors, latency is much more pronounced on a MANET relative to a traditional wired network.
- **Bandwidth:** MANET data rates are lower than dedicated enterprise links. 802.11g has a maximum rate of 54 Mbit/s vs the 1 Gbit/s common in many facilities. In addition, a less powerful signal between two nodes—often caused by factors like attenuation and interference—means lower bandwidth.
- **Connectivity:** Since the nodes in a MANET are mobile by definition, the state of the network is constantly in flux. Two nodes that had a strong direct connection may be completely unable to communicate a couple minutes later.
- **Heterogeneity:** Host computing capacity may vary widely, e.g. from supercomputers to extremely limited sensor nodes. The capacity of individual hosts may also change significantly over time: a sensor may be fully occupied while monitoring an event, but have resources to spare afterwards.

Application and coordination algorithm design must take these properties into account. Messaging primitives and the software that uses them must account for and accept latency, low bandwidth and message loss. Scarce network and device resources must also be used carefully, in addition to degrading gracefully as nodes fail or are not present, and place few assumptions on the resources that will exist. Each of these has strong implications for the messaging patterns, protocols, algorithms, and applications that may be developed on top of such networks.

## 4. DISTRIBUTED CONSTRAINT OPTIMIZATION

A constraint-based perspective can be used to formally model many essential aspects of the problem space. At the lower levels, allocation of communications spectrum, bandwidth, power, and network flow quality can all be represented as sets of constraints. At upper levels, evacuation times, triage priorities, capacities, and other problems can also be modeled as constraint optimization problems.

Of particular interest is the distributed constraint optimization (DCOP) problem model [15], in which a group of agents must choose values in a distributed fashion for a set of variables such that the cost of a set of constraints over the variables is either minimized or maximized. DCOP has generated much interest in the artificial intelligence and constraint programming communities, and a number of algorithms have been developed. The most prominent are Adopt [7], DPOP [9], NCBB [1], and OptAPO [5]. They are differentiated by the degree to which they maintain distribution of the problem and the extent to which they balance local computation time versus local memory usage.

We refer the reader to [7] for a formal definition of DCOP. Informally, the four main components of a DCOP are variables, domains, agents and constraints. Each *agent* has a set of *variables*, to which it must assign *values*. Each variable has an associated *domain*, which is the set of possible value assignments to the variable. *Constraints* are a set of functions that specify the cost of any set of partial variable assignments. Finally, each agent is assigned one or more variables for which it is responsible for value assignment. DCOP algorithms work by exchanging messages between agents, attempting to determine a globally optimal solution.

## 5. EVACUATION FORMALIZATION

A DCOP framework can model a large class of multiagent coordination and resource allocation problems can, including many in crisis management. Previous work [11, 14] has shown disparate first responder groups can coordinate to allocate tasks and develop schedules in a hostage situation. Here we demonstrate an application approach to the disaster evacuation scenario described above.

The core component of the evacuation task is assigning groups of evacuees to shelters based on sizes and requirements—matching capacities and capabilities. In an adequately provisioned scenario, no shelter is filled beyond capacity and each group's medical requirements are met. When available shelters cannot accommodate all groups in this way, the goal is to minimize some function combining costs for overflowing shelters and not meeting requirements.

Within the DCOP framework, each group is represented by an agent managing variables reflecting shelter allocations. This is a distributed analogue of the classical knapsack problem. Let $I$ be the set of items, representing groups, and $K$ be the set of knapsacks, representing shelters. The mapping $s : I \to \mathbb{N}$ relates groups to their size, and $c : K \to \mathbb{N}$ maps shelters to capacities. To encode as a DCOP, for each $i \in I$ create one variable $v_i \in V$ with associated domain $D_i = K$. Then for all possible contexts $t$:

$$f(t) \mapsto \sum_{k \in K} \begin{cases} 0 & r(t,k) \leq c(k), \\ r(t,k) - c(k) & \text{otherwise}, \end{cases}$$

where $r(t,k)$ is a function such that

$$r(t,k) = \sum_{v_i \in t^{-1}(k)} s(i).$$

The above DCOP problem formulation models the basic disaster evacuation problem of group sizes and shelter capacities. Other factors such as group requirements and the distance of a group to shelters are encoded into the DCOP as additional constraints. Group requirements can be either trauma or first aid requirements. Both group requirements and the distance traveled are encoded as "local" constraints, meaning that they are not explicitly shared with other evacuation group agents. The group requirements cost function $g : K \to \mathbb{B}$, maps shelters to a binary value indicating whether or not the requirement is met. The distance cost function, $g : K \to \mathbb{R}$, maps shelters to a real number indicating the distance to that shelter from the group's location.

Solving the DCOP ensures that all groups are assigned to shelters such that overflow is minimized, groups receive the services they need, and they travel a minimal distance.

We constructed an initial prototype of this system of tablets and PDAs using NRL OLSR[1] for ad-hoc routing support and the DCOPolis framework [12] as the distributed constraint optimization engine. DCOPolis is meant to be used for simulation as well as live testing for evaluation of DCOP algorithms under varying platforms and networking environments. It includes implementations of several algorithms and has the by-product of rendering applications independent from particular DCOP algorithms.

## 6. FUTURE DIRECTIONS FOR DCOP

Distributed constraint optimization is uniquely suited to meet the requirements imposed of coordination mechanisms in these settings. In addition to a powerful language, formalism, and algorithms for expressing and reasoning about coordination problems, the approach has several features which match well with the networking characteristics presented by disaster environments.

Most obviously, the distributed, peer-to-peer style of DCOP does not require the availability of known servers or a priori distinguished nodes, either of which cannot be guaranteed to exist. Further, distributed solving of the constraints makes use of the pooled processing and memory resources available across the network, rather than taxing any one node. This may enable networks of very limited devices to solve large, complex coordination problems.

In addition, the DCOP approach presents additional benefits, such as privacy maintenance. With some algorithms, the evaluation criterion for constraints is known only by the constrained variables' responsible agent $\alpha(v_i)$. This provides a powerful mechanism for withholding confidential information without sacrificing the global solution. It also reduces the amount of information which must be propagated around the network, reducing capacity consumption. At least one DCOP algorithm is also tolerant to message loss, delay, and asynchronous delivery.

The four most common algorithms in the literature are Adopt [7], DPOP [9], NCBB [1], and OptAPO [5]. Although all three algorithms have been proven optimal, correct, and complete they still exhibit different behavior in terms of the

[1] `http://cs.itd.nrl.navy.mil/work/olsr/index.php`

number and size of messages sent, resilience to node failure, data privacy, and synchronicity of execution.

There are several important future directions for DCOP algorithms to facilitate real world application in disaster management. The first is making the algorithms "smarter" about networking. Communication patterns of most algorithms do not incorporate network topology and routing. Utilizing network level knowledge in the solving algorithm may enable reductions in network traffic and delay. Additionally, techniques from delay tolerant networking and other networking approaches must be incorporated into solving algorithms in order to operate in these kinds of networks.

Furthermore, almost all DCOP solution algorithms assume a globally-agreed-upon tree-based ordering of the variables in the problem; distributed methods for creating these trees exist, however, the quality of the resulting trees is not guaranteed optimal which can have an adverse effect on performance. Also, the resulting trees dictate the communications patterns of the algorithms: since the trees may not coincide with the actual network topology routing becomes necessary, which again adds to the network overhead.

The next is developing algorithms that support a dynamically changing world, adapting or re-solving as the world state evolves. General techniques are needed for determining when a solution is no longer effective, dealing with the addition/loss of computation nodes, and the development of new solutions without repeating the entire solving process.

Finally, researchers need better techniques for measuring system effectiveness and performance. Several techniques have been developed for measuring the performance of a DCOP system [2, 6]. To evaluate potential solutions, evaluators are required to instrument code and run difficult to repeat experiments with each algorithm. Additionally, these metrics have many drawbacks including equations that use coefficients whose values are known to vary [4] based on four factors: the problem, the algorithm, the networking environment, and the computing power available to each agent. It is also been suggested that these coefficients are also dependent on the number of agents [10].

# 7. CONCLUSIONS

Timely and effective coordination of resources is the fundamental problem that needs to be addressed, in real-time, by personnel responding to natural disasters and other emergencies. The ability to obtain complete situation awareness, distribute tasks, and coordinate assets is the greatest obstacle toward successful response. Future systems-of-systems for these personnel will be based on a variety of emerging wireless communications technologies, none of which are perfect or infallible. Hence, coordination technologies for these future systems must be robust, decentralized, and disruption tolerant. Coordination must also be applied to domain tasks, allocating resources and effort appropriately to meet the challenges at hand.

To this end, we have presented an application of distributed constraint optimization to developing tools for disaster response and evacuation management. Requirements of this application such as decentralization and coordination make it a natural fit for such an approach. We have developed a notional coordination application and introduced an operational scenario for evaluation of the effectiveness and performance of coordination techniques in communications challenged settings. It is our belief that these techniques can be the basis of software that helps improve operational capabilities for personnel in emergency settings.

# 8. REFERENCES


[1] A. Chechetka and K. Sycara. No-commitment branch and bound search for distributed constraint optimization. In *AAMAS*, pages 1427–1429, 2006.

[2] J. Davin and P. J. Modi. Impact of problem centralization in distributed constraint optimization algorithms. In *AAMAS*, pages 1057–1063, 2005.

[3] J. Kopena, G. Naik, M. Peysakhov, E. Sultanik, W. Regli, and M. Kam. Service-based computing for agents on disruption and delay prone networks. In *AAMAS*, July 2005.

[4] R. N. Lass, E. A. Sultanik, P. J. Modi, and W. C. Regli. Evaluation of CBR on live networks. In *DCR*, September 2007.

[5] R. Mailler and V. Lesser. Solving distributed constraint optimization problems using cooperative mediation. In *AAMAS*, pages 438–445, Washington, DC, USA, 2004. IEEE Computer Society.

[6] A. Meisels, E. Kaplansky, I. Razgon, and R. Zivan. Comparing performance of distributed constraints processing algorithms. In *DCR*, 2002.

[7] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. In *AAMAS*, pages 161–168, 2003.

[8] H. Muñoz-Avila, D. Aha, L. Breslow, and D. Nau. HICAP: An interactive case-based planning architecture and its application to noncombatant evacuation operations. Technical report, Navy Center for Applied Research in AI, Naval Research Laboratory, January 1999. NCARAI Technical Note AIC-99-002.

[9] A. Petcu and B. Faltings. A distributed, complete method for multi-agent constraint optimization. In *DCR*, Toronto, Canada, 2004.

[10] M. C. Silaghi and M. Yokoo. Discussion on the three backjumping schemes existing in ADOPT-ng. In *DCR*, January 2007.

[11] E. Sultanik, P. J. Modi, and W. C. Regli. On modeling multiagent task scheduling as a distributed constraint optimization problem. In *IJCAI*, pages 1531–1536, 2007.

[12] E. A. Sultanik, R. N. Lass, and W. C. Regli. DCOPolis: A framework for simulating and deploying distributed constraint optimization algorithms. In *DCR*, September 2007.

[13] Y.-C. Tseng, M.-S. Pan, and Y.-Y. Tsai. Wireless sensor navigation for emergency navigation. *IEEE Computer*, 39(7):55–62, July 2006.

[14] T. Wagner, J. Phelps, V. Guralnik, and R. VanRiper. COORDINATORS: Coordination Managers for First Responders. In *AAMAS*, pages 1140–1147, 2004.

[15] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. The distributed constraint satisfaction problem: Formalization and algorithms. *Knowledge and Data Engineering*, 10(5):673–685, 1998.