A Multiagent System for Coordinating Ambulances for Emergency Medical Services

Beatriz López, Bianca Innocenti, and Dídac Busquets, *University of Girona*

This system uses an auction mechanism based on trust to select an ambulance for emergency patient transportation.

Emergency transportation on specialized vehicles is needed when a person's health is in risk of irreparable damage. A patient can't benefit from sophisticated medical treatments and technologies if she or he isn't placed in a proper healthcare center with the appropriate medical team. For example, strokes are neurological emergencies involving a limited amount of time in which treatment measures are effective.[1] These situations also require an expert medical practitioner to determine whether the stroke is not hemorrhagic. Only in such a case should the patient receive the medication rt-PA (recombinant tissue plasminogen activator).

Emergency medical transportation is guided by the criteria and protocols provided by regulatory authorities—for example, Spain's Emergency Medical Service (SEM). According to the SEM criteria, patients receive a transportation priority, with 0 being the highest. Each priority level requires an ambulance to arrive at the patient's location within a particular response time. For priority 0, 15 minutes is the maximum allowed time. In Spain, emergency medical transportation involves two types of ambulances. Priority-0 patients should be moved by an advanced vital support vehicle (SVA), which has an intensive-care unit. However, if no SVA is available within the required response time, a basic vital support vehicle (SVB), which has a standard care unit, can also be used. Lower-priority patients are usually moved by an SVB.

Recent studies in Girona, a rural region in northeast Spain, have shown that ambulance teams have different response times, mainly because of the drivers' expertise. In most cases, ambulances are based in a village some distance from the patient. So, the driver's knowledge of local roads is essential for arriving on time at the patient's location.

Girona has seven fully equipped ambulances (that is, SVAs with a doctor on the crew) for 223 locations spread over 5,910 km^2. The challenge for the regional authorities is to appropriately coordinate their resources to continuously improve response time, taking into account the resources needed and driver expertise. Given such a situation, the administrations of the major hospitals in Girona have considered developing a computerized system to support coordination of ambulance

1541-1672/08/$25.00 © 2008 IEEE
Published by the IEEE Computer Society

services. In response, we developed a multiagent system. Results from simulated experiments based on real data show that our system can substantially decrease response time, thus improving patient care.

The architecture

To ensure that resources (ambulances) are properly used and patients are assisted as quickly as possible, we must coordinate the different agents involved in the process. Our multiagent system provides the computational framework for such coordination.[2] The system architecture comprises ambulance team agents and an ambulance coordinator agent (see Figure 1). This model represents the current real ambulance organization for any given region. Thus, the multiagent approach lets us maintain the real distributed organization of resources.

The ambulance coordinator agent collects requests for services from other external agents (including human operators). A *service* consists of arriving at a patient's location, giving the patient first aid, and transporting the patient to the appropriate medical center. The system processes one priority-0 request at a time, as in the real organization. For each request, the ambulance coordinator agent assigns the service to an appropriate ambulance.

To solve the assignment problem, we use an auction mechanism, a well-known market mechanism to distribute tasks (in our case, services) to different agents. More specifically, we use inverse auctions (also called the contract net protocol), in which the auctioneer proposes some tasks to be performed under certain conditions.[3] The ambulance coordinator agent plays the role of auctioneer by proposing a service (one at a time) to the ambulance team agents (the bidders). The ambulance team agents reply to the ambulance coordinator with a bid containing the estimated arrival time. Using a *winner determination algorithm*, the ambulance coordinator agent selects the ambulance to which it will assign the service. The system suggests this decision to the human coordinator. If the human coordinator accepts the suggestion, the system informs the ambulance team agents and the external agent requesting the service.

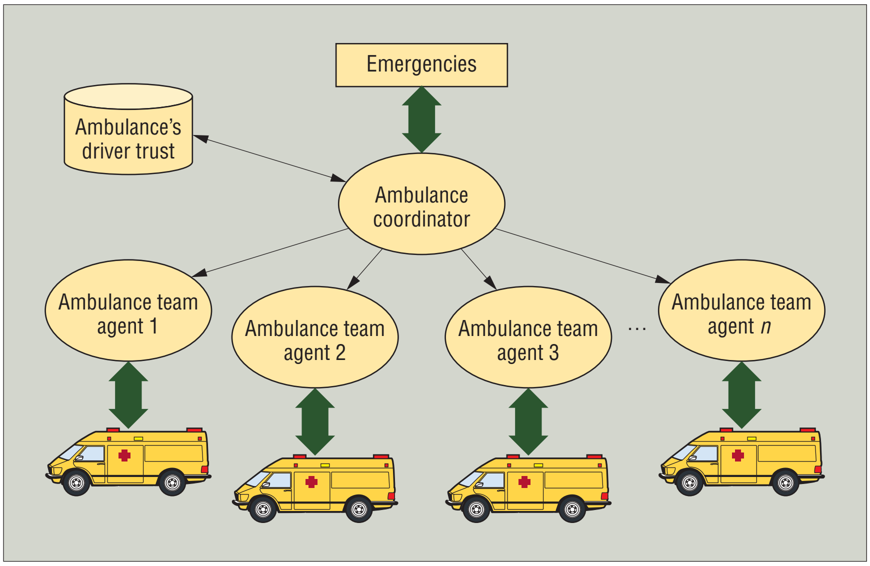

Figure 1. A multiagent-system architecture for assigning ambulances to patients. This model represents the current real ambulance organization for a given region.

Ambulance team agents

An ambulance team agent's goal is to estimate the time required to perform a service according to the ambulance's current location, availability, and crew, and the traffic conditions.

For such purposes, ambulance team agents incorporate a GPS module, a traffic module, and a trajectory module. We designed these modules in accordance with the new ambulance fleet deployed in Girona, which have been equipped with GPS and an onboard data terminal that can communicate information to and from the ambulance coordinator's PC.

The GPS module receives inputs from an ambulance's GPS system, which returns the ambulance's location. The traffic module reads information about temporarily closed roads from the national traffic center. (Our experiments used simulated GPS inputs and traffic center information.) You might think that traffic jams highly affect an ambulance's movement. However, emergency-transportation officials have told us that warning lights and sirens enable ambulances to overcome such congestion. What's really important is to avoid roads blocked by, for example, snow or roadwork. These situations would cause rerouting of the ambulance route, increasing the time to get to the patient.

The trajectory module contains a representation of the region's transportation network. Using information from the GPS and traffic modules, it calculates the optimal trajectory from the ambulance's current location to the patient's location, from that point to the hospital, and from the hospital to the ambulance base.

The system proposes the trajectory module's results to the current ambulance driver (a member of the crew). The driver can follow the proposed trajectory or adapt it according to her or his experience and knowledge about the zone where the ambulance is. When the driver accepts the trajectory, the ambulance team agent computes the estimated arrival time as a function of the distance and the maximum allowed speed (which is usually 90 kmh) for the three segments of the trajectory. If the driver rejects the trajectory, the ambulance team agent asks her or him for the estimated arrival time to the three locations.

Once the ambulance team agent has the estimated time to arrive at the patient's location (ET), to transport the patient to the hospital (TT), and to return to the ambulance's base (RT), it sends its bid to the ambulance coordinator agent. So, the bid b_i generated by the ambulance agent i has this structure:

$$b_i = \langle ET_i, TT_i, RT_i \rangle$$

The total time the ambulance requires from the moment the service is assigned until it's available again depends on not only ET, TT, and RT but also the time for supplying first aid to the patient and for hospital admission (the crew must be with the patient until she or he is attended by the hospital staff). However, the bid doesn't take into account these treatment and admission times because they don't depend on the ambulance assigned to the service.

Table 1. Ambulance coordinator experiences of ambulance team 1.

Date	Estimated time	Real time	Outcome
2 Mar.	12:00	12:15	Failure
2 Mar.	20:00	20:00	Success
3 Mar.	09:00	09:01	Success
4 Mar.	05:05	05:04	Success
4 Mar.	08:09	08:20	Failure

Obviously, an ambulance agent sends a bid only if the ambulance is available—that is, it's idle at its base or returning from a previous service. Conversely, the ambulance is unavailable from the instant it starts a service until the patient is admitted to the hospital.

The ambulance coordinator agent

As we mentioned before, the ambulance coordinator agent gathers the ambulance agents' bids and determines which ambulance should be assigned a given service. However, it would be unrealistic to assume that the estimations in the bids are always precise. To deal with possible mistakes in the estimates, we could develop a model of the driver's abilities. However, this task is difficult. Instead, we use a mechanism that approximates drivers' abilities on the basis of the ambulance coordinator agent's trust in the drivers.

The trust model

The trust an ambulance coordinator agent (a truster agent) has in an ambulance team agent is its belief that the agent can fulfill its obligations during an interaction.[4] An ambulance team agent with an expert driver should have a high trust value because that ambulance is expected to arrive on time at the patient location. Conversely, a novice driver could have a low trust value because she or he probably is less able to find alternative faster routes to the patient's location. By "expert" and "novice," we refer not to the driver's actual driving abilities but to his or her knowledge of the area covered by the ambulance.

We use Jigar Patel and his colleagues' probabilistic approach to trust. They define trust as a value in the [0, 1] interval; 0 means a totally untrustworthy agent, whereas 1 means complete reliability.[4] Because there's insufficient information to define such a probability, Patel and his colleagues propose using the expected value given previous experience of all interaction outcomes. Then, the trust value $t_{i,j}$ that agent a_i has for agent a_j is determined by

$$t_{i,j} = \frac{\alpha}{\alpha + \beta} \qquad (1)$$

where

$$\begin{aligned} \alpha &= s_{i,j} + 1 \\ \beta &= u_{i,j} + 1 \end{aligned} \qquad (2)$$

and $s_{i,j}$ and $u_{i,j}$ are the number of past successful and unsuccessful interactions.[4]

In our problem domain, a successful interaction is one in which the time the ambulance takes to arrive at the patient's location is less than or equal to the estimated time this agent sent as a bid (plus a δ margin that has been established experimentally); otherwise, the interaction fails. The initial values of α and β are 1.0.

Using this approach, we have a single truster agent, the ambulance coordinator, so we can skip the first index in the notation. The ambulance coordinator has a trust value for every ambulance team, $<t_1, \ldots, t_n>$, where n is the number of ambulance team agents. We assume that each vehicle has a permanent crew. We'll try to relax this bold assumption in future research.

The following example illustrates how we compute trust. Table 1 shows the ambulance coordinator's past experiences with ambulance team 1, assuming $\delta = 2$ minutes. From such information, $s_1 = 3$ and $u_1 = 2$. According to Equation 2, $\alpha = 4$ and $\beta = 3$. Finally, according to Equation 1,

$$t_1 = \frac{4}{4+3} = 0.57$$

Suppose that there's a new successful interaction. In this case, trust in the ambulance team increases: $s_1 = 4$, $\alpha = 5$, and, finally, $t_1 = 0.62$. If the new interaction were unsuccessful, t_1 would decrease to 0.5.

Fuzzy filtering

The ambulance coordinator agent modifies ambulance team agents' bids according to their trust values because the ambulance with the best estimated time shouldn't necessarily be the winner. It should also have a good degree of trust. So, we use fuzzy filters to modify the information provided by the ambulance team according to the team's trustworthiness.[5]

A fuzzy filter is an inference system in which the rules have this form:

if A_1 is S_1 and … and A_n is S_n **then** F is L_1

A_i and F are fuzzy variables and S_j and L_1 are fuzzy labels. A_i are the input variables; F is the filtered variable (output).

We apply the fuzzy filters three times per bid, once for each time provided in the bids: ET_i, TT_i, and RT_i. For simplicity, we focus here on ET_i; we handle the other two cases in the same way. The input variables are ET_i and trust (T_i); the filtered variable is the increasing time (IT_i), which increments the estimated time as shown in Figure 2.

For each fuzzy variable, we define these fuzzy labels:

- *ET*: {very short, short, medium, long, very long},
- *T*: {very low, low, medium, high, very high}, and
- *IT*: {very short, short, medium, long, very long}.

The fuzzy system consists of fuzzy rules such as these:

- R1: If *ET* is short and *T* is low, then *IT* is very long.
- R10: If *ET* is short and *T* is very high, then *IT* is very short.

We follow a Sugeno approach.[6,7] The input membership functions (*ET* and *T*) are triangles, and the output membership functions (*IT*) are linear. The **and** operator and the implication method are the product, and the defuzzification method is the weighted average. Figure 3 shows the system's behavior.

For each agent, we fuzzify the numerical values expressed on the bids and trust into *ET* and *T* values correspondingly (by applying the triangle membership functions). After applying the fuzzy filter, we obtain *IT*. The defuzzification of *IT* results in a numerical time increase, which modifies the original estimated time that each agent provided:

$$ET_i' = ET_i + \text{defuzzification}\left(IT_i\right)$$

Finally, we obtain a new set of estimated times for each agent:

$$\left\{ET_1', ET_2', \ldots, ET_n'\right\}$$

The winner determination algorithm

This algorithm uses the values of the modified bids to assign a task to an ambulance team.

The criterion for selecting the winner is quite straightforward: the best ambulance is the one that gets to the patient the earliest. However, we've added a small modification to take into account that not all the ambulances are identical, that the first 15 minutes after the emergency has occurred are the most critical, and that medical attention should be provided immediately.

For an emergency, an SVA is preferred. However, if an SVA will take more than 15 minutes to arrive at the patient's location but an SVB can arrive within 15 minutes, then the SVB ambulance should be selected. If no ambulance can arrive on time, the one with the shortest arrival time should be selected.

Formally, given the set A of SVAs and the set B of SVBs, with ET_i' being the estimated time of arrival of ambulance i at the patient's location (after being modified by the fuzzy filters), the selection algorithm is

$$
\begin{aligned}
&\textbf{if } \exists a \in A \text{ such that } ET_a' \le 15' \textbf{ then}\\
&\quad \textit{selected} = \arg\min_{a \in A} ET_a'\\
&\textbf{else if } \exists b \in B \text{ such that } ET_b' \le 15' \textbf{ then}\\
&\quad \textit{selected} = \arg\min_{b \in B} ET_b'\\
&\textbf{else}\\
&\quad \textit{selected} = \arg\min_{i \in A \cup B} ET_i'
\end{aligned}
$$

Maintaining region coverage

Although the winner determination algorithm ensures that an ambulance will arrive as quickly as possible, it has a dangerous drawback. It takes into account only the estimated time of arrival to the patient's location, ignoring that the ambulance must get to a hospital and, after dropping off the patient, return to its base. Thus, while an ambulance is attending a service, the area surrounding its base is undercovered.

Imagine the situation in Figure 4 (see p. 54), where a patient is located at L_p and two ambulances are available: A_1 with its base in city C_1 and A_2 with its base in C_2. Moreover, because of the treatment the patient requires, she or he must be moved to a hospital in C_2. The figure also shows the estimated driving times between each location.

According to the algorithm, A_1 should be assigned the service because its ET' is 8 minutes, whereas the ET' for A_2 is 10 minutes. However, A_1 will be away from its base for 33 minutes but A_2 will be away for only 20 minutes. So, if A_2 were assigned the

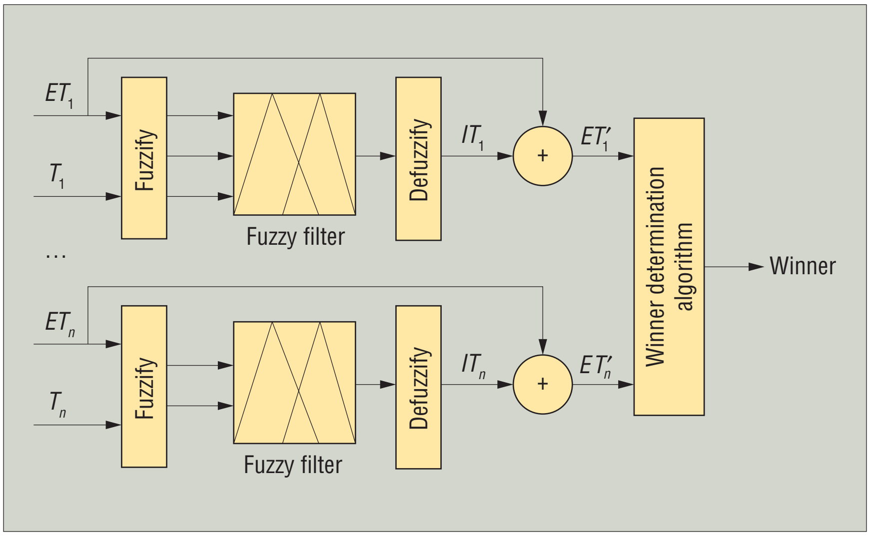

Figure 2. Fuzzy filtering. This process filters the estimated time to arrive at the patient's location (*ET*) to obtain an increasing time (*IT*), which the winner determination algorithm then uses.

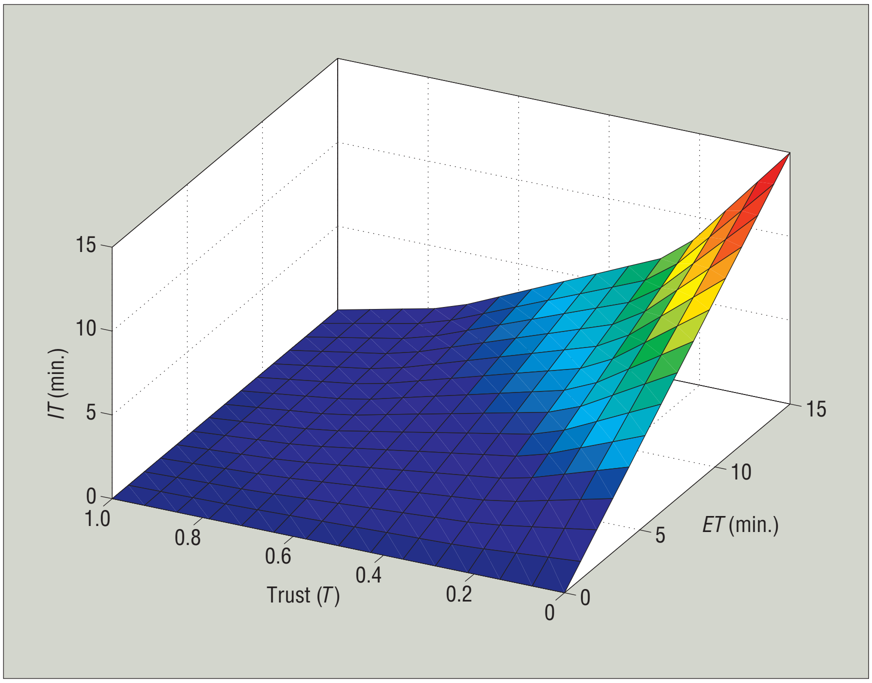

Figure 3. The fuzzy filter's output surface for any input. The figure shows the limits and the variation of increasing time (*IT*) according to the estimated time (*ET*) and the trust (*T*).

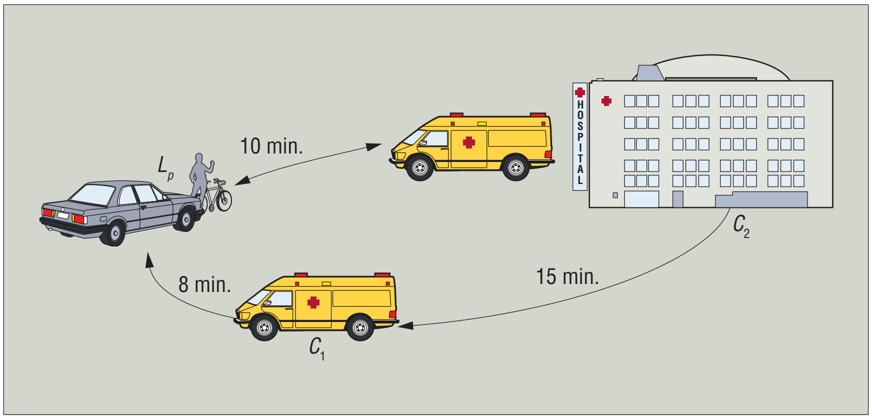

Figure 4. Driving times between the patient's location (L_p), two ambulance bases in cities C_1 and C_2, and a hospital in C_2. Although the ambulance in C_1 would arrive at the patient faster than the ambulance in C_2, its total time away from its base would be greater.

patient, she or he would still be attended within the 15-minute window, and better area coverage would be maintained (C_1 would have A_1 available, and C_2 would be unattended for only 20 minutes).

With these new criteria, we use the following *region coverage algorithm* to select one ambulance:

$$
\begin{aligned}
&\textbf{if } \exists a \in A \text{ such that } ET'_a \le 15' \textbf{ then}\\
&\quad selA = \{a \in A \mid ET'_a \le 15'\}\\
&\quad selected = \arg\min_{a \in selA}(TT'_a + RT'_a)\\
&\textbf{else if } \exists b \in B \text{ such that } ET'_b \le 15' \textbf{ then}\\
&\quad selB = \{b \in B \mid ET'_b \le 15'\}\\
&\quad selected = \arg\min_{b \in selB}(TT'_b + RT'_b)\\
&\textbf{else}\\
&\quad selected = \arg\min_{i \in A \cup B} ET'_i
\end{aligned}
$$

As in the previous algorithm, if no ambulance can reach the patient in 15 minutes, the fastest ambulance should be assigned the service.

You might think that transportation time (TT'_i) seems crucial and therefore we should try to minimize it. However, we focus only on a quick response to the emergency (ET'_i). This is because, once the ambulance reaches the given location, the medical staff can start treating and stabilizing the patient. Also, the severity of the patient's condition could cause the ambulance to stop on the way to the hospital or go very slowly to keep the patient alive. So, establishing a maximum transportation time and modifying trust in the ambulance agent on the basis of this time would be meaningless.

Experimentation

We implemented our system with three different coordination mechanisms:

- the simple assignment algorithm without the trust model,
- the simple assignment algorithm with the trust model (that is, the winner determination algorithm), and
- the region coverage algorithm with the trust model.

To evaluate our system, we used data from the regional health transportation service containing information about priority-0 ambulance services performed during nine months (approximately 7,500 services). For each service, we had

- the date and time of the call,
- the patient's location,
- the destination hospital,
- the ambulance's activation time (when it started toward the patient's location),
- the arrival time at the patient's location,
- the departure time for the hospital,
- the arrival time at the hospital, and
- the end-of-service time.

We used these data to simulate emergency calls that feed our assignment algorithm. Each service request contained the time of the call, the patient's location, and the destination hospital.

To assess our system's performance, we analyzed

- arrivals within 15 minutes,
- arrivals earlier than in the real service (we explain this later),
- the number of SVBs used, and
- the time each ambulance spent performing services and the number of services it performed.

Because information on the real ambulance assignments was missing (owing to confidentiality issues, the data didn't include the ambulance identifier), we couldn't directly compare whether our assignments were better than the real ones. However, we could indirectly compare them by comparing our ambulances' arrival time at the patient's location with the corresponding time provided in the real data. This comparison isn't strict, so we allowed a range (controlled by δ) within which times were considered the same.

Setup

We used the same number of ambulances as in the real system: seven agents for SVAs and 37 agents for SVBs. In each experiment, we initially set each ambulance agent's trust as medium (that is, 0.5). During simulation, we computed *ET* using the distance from the ambulance base to the patient's location (including the time needed for the ambulance to return to its base if it was out when it received the request). We similarly computed *TT* and *RT* using the respective origin and destination locations (the patient's location, destination hospital, and ambulance base).

However, to simulate the actual time from one location to another, we added noise to the computed values for *ET*, *TT*, and *RT*. The distribution of this noise was a function of the trust in that ambulance agent. We generated a random number between 0 and 1 and compared it with the agent's trust. If it was below the trust value, we didn't modify the driving time; otherwise, we increased it. The increment depended on the trust. For instance, for the modified *ET*, if the trust was low, the increment ranged from 0 to $|ET - \mu| + 2 \times \sigma$ min., where μ and σ are the average and standard deviation of the arrival times (simulating a driver's poor ability to avoid unexpected traffic disturbances). For medium trust, the increment ranged from 0 to $|ET - \mu| + \sigma$ min. For high trust, it ranged from 0 to $|ET - \mu|$ min. (simulating good knowledge of the road network to find shortcuts, alternative routes, and so on). With these ranges, an agent with low trust can increase its trust and an agent with high trust can still be penalized.

Table 2. Simulation results (average ± deviation). All figures are percentages.

	Without trust	With trust	With trust & region coverage	Real data
Arrivals within 15 min.	97.4340 ± 0.2217	99.8459 ± 0.0559	99.5311 ± 0.0565	87.57
Early arrivals	89.4252 ± 0.1652	92.6986 ± 0.1558	93.8706 ± 0.2579	—
SVB ambulances used	12.2632 ± 0.0928	13.9569 ± 0.0857	8.4874 ± 4.8701	—

We then compared this simulated driving time against the estimated time sent by the agent, to update its trust. The value of the parameter δ was 2 minutes. in all the experiments. We took the times spent on the scene to treat the patient and in hospital admission from the real data of each service.

Currently, we use no realistic simulator that keeps track of the ambulances' locations at any given time. This affects our simulation in two ways. First, the estimated driving times aren't very precise. Second, assume an ambulance is in the middle of a service and receives an assignment for another service. To calculate the time to get to the patient, the ambulance team agent computes the time the ambulance takes to return to its base plus the time from the base to the patient. However, in some such real cases, the ambulance could divert its route in advance, thus decreasing the time to get to the patient.

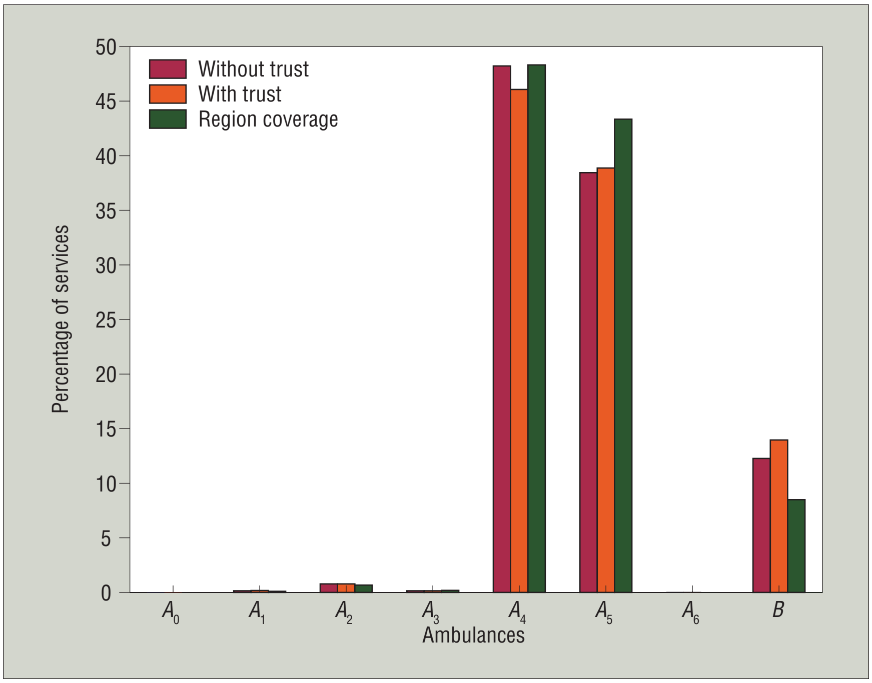

Figure 5. The distribution of services among ambulances. The region coverage algorithm significantly reduces the number of SVBs (basic vital support vehicles).

Results

Table 2 shows the results of our simulations for each coordination mechanism, together with the real data. The results are averages over 10 different simulations with the same real data. We performed Student's t-test, and all the results are statistically significant.

Almost all the services were performed within 15 minutes or less. This percentage (above 97 percent in all the experiments) is higher than that achieved in the real data, where only 87.57 percent of the services were attended within 15 minutes. Moreover, the small percentage of services that failed to arrive within the window in our experiments were cases in which the ambulance arrived at the patient's location after 16 to 18 minutes; no case suffered a significant delay.

With the simple assignment algorithm, approximately 90 percent of the services had performance times as good as or better than the real services' times. The introduction of trust slightly improved this percentage, increasing it by 2.4 percentage points. The region coverage algorithm significantly reduced the number of SVBs used. This indicates that SVA assignment improved and that SVAs could therefore provide more services without needing to request an SVB.

Reduced SVB use is also evident in Figure 5, which shows the percentage of services performed by each of the seven SVAs and the total for all the SVBs. Ambulances A_4 and A_5 performed most of the services (approximately 90 percent). This is because both ambulances were in the city of Girona (the largest in the region) and most of the services were in this city. With the region coverage algorithm, these two ambulances performed a higher percentage of services, whereas the percentage for the SVBs decreased.

However, ambulances A_0 and A_6 were seldom (if ever) used because they were in low-population areas. Although having these two ambulances idle most of the time could seem a waste of resources, it wasn't. These areas are far from big cities; if they didn't have an ambulance, it would take a long time for another ambulance to reach them.

This uneven distribution of services also prevented us from observing a greater effect of including trust in the assignment process. This is because, given that most services were in the city of Girona, the two ambulance teams there were usually the best, regardless of their trust.

To see the region coverage algorithm's impact, we also compared the time each ambulance was busy. Figure 6 (see p. 56) shows, for each SVA, the simulated driving times according to the random distribution based on trust. "Going with trust" is the time an ambulance took from its base to the patient when the system used trust, and "Returning with trust" is the total time to return to the base, again when using trust. "Going with region coverage and trust" and "Returning with region coverage and trust" are the same total times when the system used region coverage and trust. Although the region

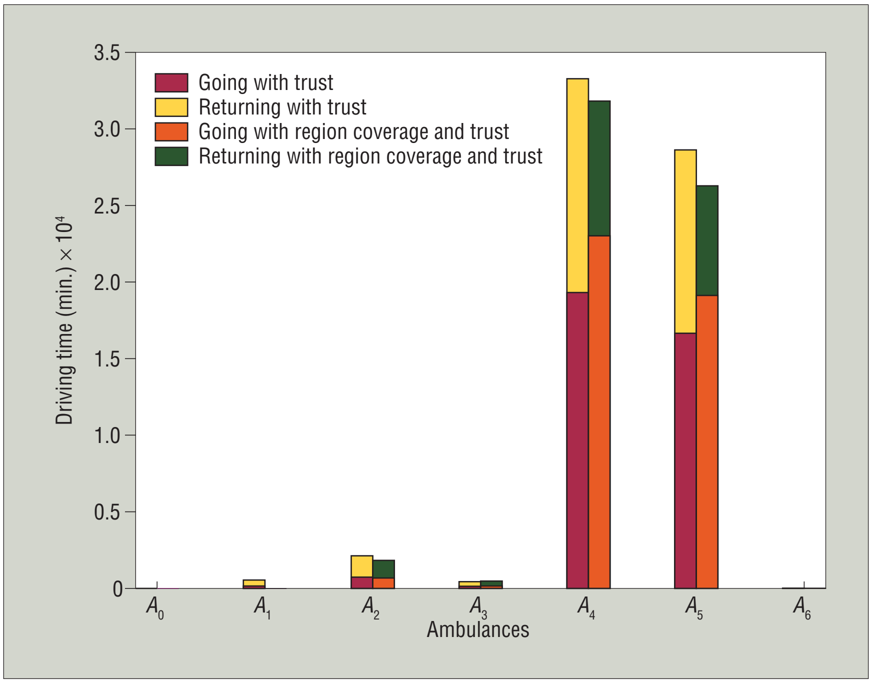

Figure 6. The driving time for each ambulance. Although with region coverage, the time to get the patients increases, each ambulance's total time in service decreases.

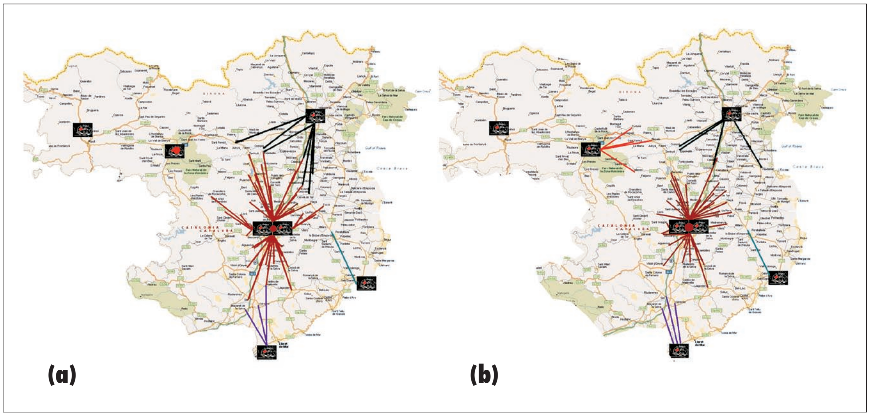

Figure 7. The assignment of services to ambulances, taking into account (a) region coverage and trust and (b) just trust. A comparison of both figures shows that with the region coverage algorithm, the ambulances based in Girona cover a larger area owing to their proximity to the main hospital.

coverage algorithm increased the time to go to patients (because the selected ambulance wasn't always the fastest), it reduced the return time and the total time each ambulance was on the road. This implies that the ambulances were available more often, thus giving better coverage to the region.

Figure 7 also illustrates the region coverage algorithm's impact. It shows the services each ambulance performed (the rectangles indicate individual ambulance bases, with the lines pointing to the different patient locations to which they traveled). With the region coverage algorithm (see Figure 7a), the two ambulances in the city of Girona covered an area larger than they did without such coverage (see Figure 7b). This is because these ambulances had low return times and therefore were assigned most services that the simple assignment algorithm assigned to ambulances at other bases.

Our results show that a multiagent system can integrate the components of an ambulance organization so that the system can assign the most appropriate vehicle for emergency patient transportation. (For a look at other applications of agents to emergency medical services, see the sidebar.) The system can provide a quick response, the first step to ensure that the patient receives proper treatment. By combining an auction protocol with a trust model and fuzzy filters that deal with driver expertise, the system can take into account more variables in the decision process than current assignment strategies do. In addition, maintaining region coverage improves decisions regarding ambulance distribution. Nevertheless, such a coverage strategy is still an experimental coordination model that should be discussed with medical authorities, who are ultimately responsible for deciding the trade-off between coverage and ambulance arrival time.

Our auction mechanism based on trust has room for improvement. The computation of trust is based on a strict criterion (the ambulance either arrived in time or was late). A fuzzier approach could take into account the actual difference between the estimated time and the real arrival time, better modeling the driver's expertise.

Moreover, the region coverage algorithm assumes that each base has only one ambulance, so when that ambulance is busy, the area is undercovered. However, a real base might have several ambulances, so we need to extend the algorithm to deal with this issue.

Another important goal is to integrate this multiagent system with other healthcare systems dealing with medical protocols and diagnosis,[2] which would automatically request ambulance service when needed. As part of this integration, we could extend our system to cover other ambulance services by including lower-priority (but still urgent) patient transportation and cases that aren't urgent. In the nonurgent scenario, a quick response isn't needed; the problem is to assign sets of patients to different ambulances according to the patients' transportation requirements. To solve this problem, we're following a combinatorial auction approach.[8]

Acknowledgments

The Spanish MEC (Ministerio de Educación y Ciencia) project TIN2004-06354-C02-02 and the DURSI (Departament d'Universitats, Recerca

Related Work on Agents and Emergency Medical Services

The past few years have seen more interest in research on managing disaster scenarios (such as earthquakes and hurricanes) than in handling medical transport. Scenarios such as Robocup Rescue (www.robocuprescue.org) involve coordinating other teams in addition to ambulances, such as police forces and fire brigades. Such scenarios are so complex, and each proposed solution involves so many different techniques, that evaluating a single component, such as ambulance allocation, is difficult. Wei Chen and Keith Decker tried to abstract from the disaster scenario by working only on ambulance-police coordination. However, they focused on how and when the different agents should communicate to coordinate their actions.[1] The research we present in the main article is more specific, and it addresses the problem of assigning services to ambulances.

Regarding multiagent approaches for coordinating agents involved in medical transport, Anna Ciampolini and her colleagues propose having the ambulance choose the destination hospital.[2] In contrast, in our system, the hospital is known in advance, and the system assigns the ambulance to transport the patient. Our approach is in line with other medical coordination approaches, such as those of Decker and Jinjiang Li[3] and David Isern and Antonio Moreno.[4] To some extent, the multiagent systems in these approaches replicate existing human organization and authority structures. This is also the case in our approach; we believe that using a different structure would be unacceptable to most medical authorities.

References

1. W. Chen and K.S. Decker, "The Analysis of Coordination in an Information System Application—Emergency Medical Services," *Agent-Oriented Information Systems II*, LNCS 3508, Springer 2005, pp. 36–51.
2. A. Ciampolini, P. Mello, and S. Storari, "A Multi-Agent System for Medical Services Synergy and Coordination," *Proc. 2nd Workshop Agents Applied in Health Care*, 2004, pp. 38–46.
3. K. Decker and J. Li, "Coordinating Mutually Exclusive Resources Using GPGP," *J. Autonomous Agents and Multi-Agent Systems*, vol. 3, no. 2, 2000, pp. 133–158.
4. D. Isern and A. Moreno, "Agent-Based Careflow Using CPGs," *Recent Advances in Artificial Intelligence Research and Development*, IOS Press, 2004, pp. 293–300.

i Societat de la Informació) Automation Engineering and Distributed Systems Group, 00296, partially funded this research. Many thanks to Esteve del Acebo for his fruitful discussions about fuzzy filters. We also thank Quim Llorella and Rosa Pérez for sharing details of the real emergency transportation service.

References

1. A. Davalos and M. Fisher, "Emerging Therapies for Cerebrovascular Disorders," *Stroke*, vol. 36, no. 2, 2005, pp. 208–210.
2. D. Isern and A. Moreno, "Agent-Based Careflow Using CPGs," *Recent Advances in Artificial Intelligence Research and Development*, IOS Press, 2004, pp. 293–300.
3. M.P. Wellman and P.R. Wurman, "Market-Aware Agents for a Multiagent World," *Robotics and Autonomous Systems*, vol. 24, nos. 3–4, 1998, 115–125.
4. J. Patel et al., "A Probabilistic Trust Model for Handling Inaccurate Reputation Sources," *Trust Management*, LNCS 3477, Springer, 2005, pp. 193–209.
5. G.J. Klir and T.A. Folger, *Fuzzy Sets, Uncertainty and Information*, Prentice Hall, 1992.
6. M. Sugeno, "On Stability of Fuzzy Systems Expressed by Fuzzy Rules with Singleton Consequents," *IEEE Trans. Fuzzy Systems*, vol. 7, no. 2, 1999, pp. 201–224.
7. E. del Acebo and J.L. de la Rosa, "A Fuzzy System Based Approach to Social Modeling in Multi-Agent Systems," *Proc. 1st Int'l Joint Conf. Autonomous Agents and Multiagent Systems* (AAMAS 02), part 1, ACM Press, 2002, pp. 463–464.
8. T. Sandholm, "Algorithm for Optimal Winner Determination in Combinatorial Auctions," *Artificial Intelligence*, vol. 135, nos. 1–2, 2002, pp. 1–54.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.

The Authors

Beatriz López is an associate professor in the University of Girona's Department of Electrical, Electronics and Automatic Engineering. Her research interests include multiagent systems, planning and scheduling, and case-based reasoning. López received her PhD in computer science from the Technical University of Catalonia. She's a member of the Catalan Association for Artificial Intelligence. Contact her at beatriz.lopez@udg.edu.

Bianca Innocenti is an assistant professor in the University of Girona's Department of Electrical, Electronics and Automatic Engineering. Her research interests are mobile robotics and multiagent systems. Innocenti obtained her DEA in industrial informatics and advanced control technologies from the University of Girona. Contact her at bianca.innocenti@udg.edu.

Dídac Busquets is a research scientist in the University of Girona's Department of Electrical, Electronics and Automatic Engineering, where he's applying market mechanisms to task allocation and multiagent coordination. Busquets received his PhD in computer science from the Technical University of Catalonia. Contact him at didac.busquets@udg.edu.