Modelling an Emergency Vehicle Early-Warning System using Real-time Feedback

Aline Senart*, Mélanie Bouroche and Vinny Cahill

Distributed Systems Group, Department of Computer Science
Trinity College Dublin, Ireland
E-mail: aline.senart@cs.tcd.ie
*Corresponding author

Abstract: Emergency vehicles, while usually equipped with warning sirens and/or beacons, are still often impeded by other traffic and involved in numerous collisions. An effective warning system for alerting other vehicles that an emergency vehicle is approaching is therefore crucial. Current research is investigating the use of wireless communication from emergency vehicles to warn other traffic. However, communication in dynamic wireless networks is not reliable. In this paper, we propose a novel emergency vehicle early-warning system that relies on wireless technology to inform other vehicles of the arrival of emergency vehicles, and returns feedback in real-time when communication degrades. When such a situation happens, emergency vehicles are informed that some vehicles may not have been warned and should therefore slow down to negotiate traffic safely. We describe how such a system is built with Comhordú, a coordination model for wireless networks.

Keywords: Emergency vehicle warning system, wireless ad hoc network, real-time feedback, coordination.

Reference to this paper should be made as follows: Senart A., Bouroche M. and Cahill V. 'Modelling an Emergency Vehicle Early-Warning System using Real-time Feedback', *Int. J. Intelligent Information and Database Systems*, Vol. x, No. x, pp.xxx–xxx.

Biographical Notes: Aline Senart is a Senior Research Fellow in the Distributed Systems Group at Trinity College Dublin. She received an M.Sc. in Computer Science, Systems and Communication from Joseph Fourier University in 2000 and her Ph.D. in Computer Science from INP Grenoble in 2003. Aline's Ph.D. work was on dynamically reconfigurable operating system kernels with application to active networks. Her current research interests are in middleware architectures and programming models to support mobile and context-aware applications ranging from intelligent transportation systems to emerging global supply chain management.

Mélanie Bouroche is a Ph.D. student in the Distributed Systems Group at Trinity College Dublin. She was awarded a "Diplôme d'Ingénieur en Télécommunications" by INP Grenoble, France, and an M.Sc. in Com-

Copyright © 200x Inderscience Enterprises Ltd.

puter Science (Networks and Distributed Systems) from Trinity College Dublin in 2003. Her research interests include distributed systems, and in particular, middleware for sentient computing, and real-time coordination of mobile autonomous entities. Mélanie's thesis presents an approach to systematically translating system-wide safety constraints into requirements on the behaviour of autonomous mobile entities.

Vinny Cahill holds a Personal Chair in Computer Science and is a Science Foundation Ireland investigator at Trinity College Dublin where he also serves as Director of Research for Computer Science and Statistics. His research addresses many aspects of distributed systems, in particular, middleware and programming models for ubiquitous and mobile computing with application to intelligent transportation systems and personal healthcare/independent living. He has over 100 peer-reviewed publications in international conferences and journals and was a founding editorial board member of IEEE Pervasive Computing Magazine and editorial board member of IEEE Distributed Systems Online.

1 Introduction

As traffic volumes increase on roads, emergency responders (e.g., ambulances, fire trucks and police vehicles) face a growing risk to their occupants' personal safety while approaching the scenes of accidents. While usually equipped with sirens and/or beacons, the sudden appearance of emergency vehicles can be extremely disruptive to nearby vehicles. Corporation (2002) reports that emergency vehicle crashes occurring while responding to emergencies represent one of the major causes of injuries and deaths of professional responders. Moreover, warning lights and sirens, which are traditionally used by emergency service providers when called to accidents, are time-saving (Brown et al., 2000) but leave little time for vehicles to maneuver to get out of the way. Some drivers become confused and create conflicts with other drivers that can block lanes and increase vehicle response times. However, every minute counts in emergency situations. For example, when a person goes into defibrillation, waiting for an emergency vehicle to arrive eats away precious time (Eisenberg, 1998). An effective early-warning system for alerting the drivers of other vehicles (or in the future notifying autonomous vehicles) that an emergency vehicle is approaching is therefore crucial.

Radio signaling systems can also be used to provide advanced warning of an approaching emergency vehicle, thereby potentially reducing traffic accidents by increasing the awareness of the surrounding vehicles. Such a warning system can be very effective as information about the arrival of an emergency vehicle can be relayed by the roadside infrastructure or the vehicles themselves, and this beyond the driver's horizon (Santos et al., 2004; Little and Agarwal, 2005; Farkas et al., 2006).

Current research is therefore investigating the use of wireless communication for emergency vehicles to notify traffic lights and private and public transport vehicles of their arrival (Miyawaki et al., 1999; Durresi et al., 2005). While these warning systems contribute to safer and more efficient roads, wireless communication is not reliable in dynamic vehicular networks characterised by high mobility and speed

(Yin et al., 2004). If transmissions are delayed or lost, the safety of emergency vehicles and other vehicles can be compromised. To allow vehicles to make safe progress in the presence of unreliable communication, it is necessary to inform them when timely information cannot be provided.

In this paper, we propose a novel emergency vehicle early-warning system that relies on wireless technology to inform private and public transport vehicles of the arrival of emergency vehicles and returns feedback in real-time to emergency vehicles about the maximum speed at which they can safely drive. Therefore, when real-time communication cannot be guaranteed, emergency vehicles are informed that nearby vehicles may not be warned within a suitable time bound. They then have the possibility of adapting their behaviour to ensure that road safety constraints are not violated, for example, by slowing down to negotiate traffic safely while still using traditional warning systems.

We describe how such a system is built with Comhordú (Bouroche et al., 2006a,b), a decentralised coordination model for wireless networks, in which mobile entities coordinate their behaviour to ensure that application-specific safety constraints are never violated. This coordination model builds on a real-time communication model for wireless networks, known as the space-elastic model (Hughes, 2006), that provides feedback to entities about the state of communication. With the coordination model, entities can use this feedback to ensure that the safety constraints are always satisfied, by adapting their behaviour when communication is limited.

The remainder of this paper is structured as follows. In Section 2, we review related work on emergency vehicle warning systems, autonomous vehicles and coordination models. In Section 3, we describe the space-elastic model for real-time communication in wireless networks. In Section 4, we present Comhordú, a coordination model in which safety constraints are not violated even in the face of possible communication failures, while in Section 5 we describe how the emergency vehicle early-warning system using real-time feedback is built with the coordination model, and show that warning sytems with feedback allow emergency vehicles to drive faster than without feedback when wireless communication is sufficiently good. Finally, Section 6 concludes this paper.

2 Related work

Emergency vehicles generally depend on sirens, horns, flashing lights or some other type of audible or visible alarm to alert other vehicles in the area as they approach. A few systems have been proposed or patented that use radio communication to provide additional warnings (Mayhew and Shirkey, 1992; Code et al., 2006; Ewing and Zubiel, 2006), some of which have been commercialised (NASA, 2001). Other systems use wireless communication to provide automatic remote control of traffic lights at intersections, so as to switch to a green light for the direction of the approaching emergency vehicle and to a red light for opposing traffic (Miyawaki et al., 1999; Casturi, 2000; Dötzer et al., 2005; Transportation Research Board, 2006).

Similar work can be found in the broader area of autonomous vehicles. The challenges involved in realising this vision of driverless vehicles are similar to those faced by the development of emergency vehicle warning systems in terms of the timeliness constraints on communication that are required. A number of projects have been conducted in the autonomous vehicles domain and what was considered science fiction not long ago is now becoming a reality. Prototypes have been designed that are able to overtake a slower vehicle, cross an unsignalised junction and avoid collision with other vehicles (Kolodko and Vlacic, 2003; Baber et al., 2005; Pallottino et al., 2007). A similar scenario is also presented in (Mock, 2004), where robots coordinate their actions to cross a shared road section.

There have also been numerous efforts addressing the coordination/cooperation of mobile entities. A communication architecture for the cooperation of autonomous vehicles is presented in (Nett et al., 2001) and detailed in (Schemmer et al., 2001), however their approach is limited to infrastructure-based networks. Similarly, existing coordination models for mobile autonomous entities, such as Linda In Mobile Environments (LIME) (Murphy et al., 2001) and EgoSpaces (Julien and Roman, 2007) facilitate rapid development of mobile applications. EgoSpaces demonstrated the usefulness of their middleware through the easy construction of an emergency vehicle warning system that turns on a light on the dashboard of a car when the driver needs to clear the road for an emergency vehicle.

In all this work, either reliability of communication is assumed, or unreliability is tackled by retransmission of signals until correct reception (may be infinite). In both cases, timely message delivery cannot be assumed. Therefore these approaches can only increase the probability that vehicles are warned in advance of the arrival of emergency vehicles. However, as only a best-effort service is provided, emergency vehicles do not have guarantees that vehicles are warned (some signals may be delayed or never arrive). Therefore, the speed of emergency vehicles is still limited by their line of sight as they have to be ready to stop in case some vehicle did not receive the message and free the way.

Finally, recent advances in wireless technologies provide new opportunities for advanced emergency vehicle early-warning systems. In particular, Dedicated Short Range Communication (DSRC), a short- to medium-range wireless protocol specifically designed for automotive use, offers the potential to effectively support high-speed inter-vehicle and vehicle-to-roadside communications (Xu et al., 2004). DSRC is already used in electronic toll collection, electronic credentialing and monitoring of commercial vehicle operations and could be applied in many different other scenarios including warning systems (Tsugawa, 2002; Lee et al., 2004). However, to our knowledge, no such system has been designed so far that provides the kind of feedback required by emergency vehicle warning systems.

3 Space-elastic model

The space-elastic model (Bouroche et al., 2006a; Hughes, 2006) is a model for real-time communication in wireless networks, including ad hoc networks. In this model, real-time communication is guaranteed within a geographical proximity of the sending entity. This proximity can be of any shape and can be defined either absolutely (e.g., using GPS coordinates), or relatively around the entity (using an anchor point and a size).

3.1 Specifications

An entity wishing to send messages specifies the proximity within which it wants its messages to be delivered. This proximity is called the *desired coverage*, and is used to bound message propagation. The entity also stipulates the maximum latency, $msgLatency$, within which the messages must be delivered, and the desired period for their transmission, named *period*.

Depending on the topology of the network (i.e., the distribution of the nodes and the quality of the wireless links), it might not be possible over some period of time to deliver a message in time to all interested entities within the desired coverage. Therefore, the size of the proximity in which timely delivery of messages is provided, called the *actual coverage*, changes over time. In the worst case, no communication is possible; this corresponds to an actual coverage of size zero. The sender is notified in real-time of any expansion or shrinking of the actual coverage within a bounded time, $adaptNotif$. Therefore, an entity knows within $msgLatency + adaptNotif$ after sending a message the area in which it has been successfully delivered, and can adapt its behaviour accordingly. If the actual coverage becomes smaller than one or more thresholds, called *critical coverage(s)*, the sender might need to take into account that it cannot communicate in an area wide-enough to maintain safe operation, and might need to adapt its behaviour. Possible variations of the actual coverage around the desired and the critical coverages are illustrated in Figure 1.

3.2 Guarantees

An entity is said to be present within the actual coverage of a sending entity once it is able to receive messages after arriving in communication coverage. This takes an implementation-dependent time, $present$, which is necessary to include the entity in the real-time routing protocol for example. The real-time communication guarantees that the space-elastic model provides are therefore:

- to message senders: to be able to communicate within $msgLatency$ in the actual coverage, and to be notified within $adaptNotif$ if this coverage changes,
- to message receivers: to receive every message of the types in which they have expressed interest, if they are present within the actual coverage of the message sender at the delivery time of this message.

The feasibility of the space-elastic model with low-jitter real-time communication and time bounded adaptation notification has been demonstrated in real-world settings (Hughes, 2006). We show in the next section how the feedback describing the revised actual coverage in which communication is guaranteed can be used to ensure safety. This is the basis for the design of our emergency vehicle warning system with real-time feedback.

4 Coordination model

In this section, we introduce Comhordú (Bouroche et al., 2006a,b), a real-time coordination model for mobile autonomous entities based on the notion of distributed responsibility. We first present the formalism used to express high-level system-wide safety constraints. We then show how these safety constraints are translated with Comhordú into requirements on the behaviour of individual entities.

4.1 Specifying the safety constraints

Safety constraints typically include constraints on the actions and states of entities, as well as their distance to each other. For example, a vehicle needs to coordinate its behaviour with an emergency vehicle only when they are on the same stretch of road. In this section, we introduce a formalism to define such safety constraints.

4.1.1 Scenario, modes and states

A *scenario* encompasses a set of *entity types* $E_1, E_2, .., E_n$, a *goal*, and some *safety constraints*. In the emergency vehicle warning system, the entity types represent emergency vehicles and other vehicles. The goal of the system is for emergency vehicles to drive as fast as possible to their destination, and for the other vehicles to make safe progress to their destination. The safety constraints are that the emergency vehicle should not crash into another vehicle.

The behaviour of an entity is described by a set of *modes* that describe the actions that can be taken, and the possible transitions between these modes. Transitions between modes are instantaneous. For example, given the maximum velocity of an emergency vehicle $v_{\max}$, and an increasing set of speeds $\{v_i\}_{i\in[1,n]}$ with $v_n = v_{\max}$, the modes of emergency vehicles can be defined as: `stopped`, $\{$`going_at_v`$_i$, `accelerating_to_v`$_i$, `braking_to_v`$_i\}_{i\in[1,n]}$. Similarly, the behaviour of the other vehicles can be modelled with the modes `travelling`, `getting_out_of_the_way`, and `out_of_the_way`. We denote the set of modes of entities of type E_i as M_i.

The situation of an entity at a given time is described by its *state* which encompasses its mode, and some additional application-specific information, for example, the position of the entity. We denote the set of states of entities of type E_i as S_i. The states of emergency vehicles and other vehicles encompass their mode, location, current speed and direction.

4.1.2 State compatibility

A set of states $(s_1, s_2, ..., s_n) \in \mathrm{S}_1 \times \mathrm{S}_2 ... \times \mathrm{S}_n$ is *compatible,* noted $\mathcal{C}_{\mathrm{s}}(s_1, s_2, ..., s_n)$, if the safety constraints are not violated when some entities are simultaneously in these states. For instance, the states of an emergency vehicle and other vehicles are compatible if they are far enough away. Also, the state of a vehicle that is off the road is compatible with the state of any emergency vehicle.

4.1.3 Expressing the safety constraints

The safety constraints can be expressed as a set of incompatibilities between states, including constraints on the relative distance of entities (noted distance(*position*1, *position*2)). For example, the safety constraint that emergency vehicles and other vehicles should not collide can be stated as:

$$\mathcal{C}_\mathrm{s}(s_\mathrm{OV}, s_\mathrm{ev}) \text{ iff } \neg\big((\text{distance}(s_\mathrm{OV}.\text{position}, s_\mathrm{ev}.\text{position}) < d) \wedge (s_\mathrm{ev}.\text{mode} \neq \texttt{stopped}) \wedge (s_\mathrm{OV}.\text{mode} \neq \texttt{out\_of\_the\_way})\big)\,, \quad (1)$$

where s_ev refers to the emergency vehicle, s_OV to any other vehicle, and d to a threshold distance. This equation expresses the fact that the safety constraints will not be violated if a vehicle and an emergency vehicle are far enough away, or the emergency vehicle is stopped, or the vehicle is out of the way of the emergency vehicle.

4.2 Ensuring the safety constraints

High-level system-wide safety constraints, while being simple and quite intuitive to state, are not easily exploitable as such. Comhordú introduces a number of abstractions that enables to deduce the necessary and sufficient requirements on individual entities' behaviour so that these safety constraints will not be violated.

4.2.1 Responsibility

For every state incompatibilities between entities, one of the entities needs to ensure that the incompatibility will not occur. This entity is *responsible* for the incompatibility. For example, emergency vehicles are responsible to ensure that they do not collide with other vehicles.

Responsibility might be attributed a priori or in real-time, and might be transferred. However, at any time, at least one entity must be responsible for each possible incompatibility. This notion of responsibility allows to distribute the duty of ensuring safety constraints over entities.

An entity responsible for an incompatibility needs to foresee when the incompatibility might happen in order to prevent it. Occurrences of incompatibilities can be deduced from the modes of the different entities.

4.2.2 Mode compatibility

A set of modes $(m_1, m_2, ..., m_n) \in \mathrm{M}_1 \times ... \times \mathrm{M}_n$ is *compatible* if, when some entities are simultaneously in these modes, their states are always compatible. Given $\mathrm{S}_{i,m} \subseteq \mathrm{S}_i$ a set of states of the entity e_i in which it is in mode $m \in \mathrm{M}_i$, mode compatibility of the modes $m_1, m_2, ..., m_n$ of entities $e_1, e_2, ..., e_n$ can be defined as:

$$\mathcal{C}_\mathrm{m}(m_1, m_2, ..., m_p) \text{ iff } \forall (s_1, s_2, ..., s_n) \in \mathrm{S}_{1,m_1} \times ... \times \mathrm{S}_{n,m_n},\, \mathcal{C}_\mathrm{s}(s_1, s_2, ..., s_n)\,.$$

For example, the modes `out_of_the_way` of a vehicle and $\texttt{going\_at\_v}_i$ of an emergency vehicle are compatible because when they are in these modes, their states are always compatible.

While the notion of state incompatibility captures whether the safety constraints are being violated at a given time, mode compatibility enables to predict that no incompatibility will happen if entities are in these given modes. Note that if the modes of a set of entities are not compatible, it does not imply that the safety constraints will be violated. For example, the modes `travelling` of a vehicle and `going_at_v`$_i$ of an emergency vehicle are not compatible, as entities might collide into each other when they are in these modes. However, if they are far enough apart, the safety constraints will not be violated.

A mode of an entity e_i is said to be a *fail-safe mode* if e_i can remain in this mode to ensure that the incompatibility for which it is responsible will not happen. For example, the mode `stopped` is a fail-safe mode for emergency vehicles.

4.2.3 Coordination primitives

To ensure that no state incompatibility will happen, a responsible entity can use three primitives: it can adapt its behaviour, delay its own actions, or transfer its responsibility.

Adapting its behaviour (`adapt`) A responsible entity can obtain information about the modes that other entities can be in by both previous knowledge and in real-time. Using this information, a responsible entity can adapt its behaviour, e.g., change mode, to always be in a mode that ensures that the safety constraints will never be violated.

Delaying actions (`delay`) A responsible entity can ensure that the incompatibility for which it is responsible will not happen by delaying an action that can trigger this incompatibility, e.g., delay switching to a mode in which an incompatibility might occur. The action can be delayed until either the responsible entity gets information that it is safe to undertake it, or it has warned all entities that it will undertake it.

Transferring responsibility (`transfer`) The last means that can be used by a responsible entity to ensure that the incompatibility for which it is responsible will not occur is to warn other entities that the incompatibility might occur. The responsible entity might include its state and mode in the warning message. The other entities are expected to change their behaviour to avoid the incompatibility and can use information contained in the warning message to do so. Warning messages need to be sent periodically over a proximity that is wide enough to ensure that entities approaching will be notified early enough so that they will have time to react.

An entity sending a message receives feedback about the delivery area within a time bound, but not whether there was any entity within this area that actually received the message. Therefore, entities that receive the message become responsible to ensure that no incompatibility arises with the sender, but this transfer is however only partial, as the sender remains responsible for the incompatibility in relation to other entities.

4.3 Translating the safety constraints

In this section, we introduce contracts and zones, which build on responsibility, mode compatibility and the coordination primitives to allow the translation of the safety constraints into requirements on the behaviour of entities.

4.3.1 Contracts between entities

Three types of contracts between entities exist and we detail each of them below.

Contract without transfer The responsible entity does not transfer its responsibility, and must always ensure that the safety constraints are not violated, by adapting its behaviour if necessary. Other entities do not need to be aware of the contract. Therefore, the only coordination primitives that can be used are `adapt` and `delay` and are used only by the responsible entity.

Contract without feedback The responsible entity must warn other entities at least t_{warning} in advance when the safety constraints are liable to be violated. Other entities must be able to react within this t_{warning} to ensure that no incompatibility will happen. For this contract, the responsible entity can use any coordination primitive, while other entities are restricted to `adapt`.

Contract with feedback In this contract, the responsible entity must also warn other entities at least t_{warning} in advance when the safety constraints are liable to be violated, but entities can in addition provide feedback to the responsible entity when they cannot adapt their behaviour. In turn, the responsible entity must also be able to react to ensure that no incompatibility will happen, and this within $t_{\text{warning}} - t_{\text{feedback}}$. Other entities must be able at any time either to react within t_{warning} to a message from a responsible entity, or to communicate within t_{feedback} to this entity. In this contract, both responsible and other entities can use the three primitives.

The use of the three primitives by both responsible entities (R) and others (O) in the three contracts is described in Table 1.

Table 1 Use of the primitives by the contracts

Contract	`adapt`	`delay`	`transfer`
Without transfer	R	R	-
Without feedback	R, O	R	R
With feedback	R, O	R, O	R, O

In the emergency vehicle warning system, if emergency vehicles are initially responsible for avoiding collisions with other vehicles, the contract without feedback between emergency vehicles and other vehicles could be used: emergency vehicles have to warn other vehicles early enough before their arrival, so that other vehicles get out of their way.

4.3.2 Zones

The contracts between entities can be translated into geographical zones: the safety and consistency zones, and the critical coverage.

Safety zone The states of all the entities must be compatible at all times, but the safety constraints actually impose constraints only on specific states, typically when two entities are "close". For this reason, the *safety zone*, denoted SZ, is defined as the set of positions of entities in which their states are liable to be incompatible with those of the responsible entity.

Consistency zone If a responsible entity foresees that an entity could be in a state that is not compatible with its own state when that entity enters its SZ, the responsible entity can choose to transfer its responsibility. In this case, it must do so early enough, so that the incoming entity have time to adapt its behaviour (either by not entering the safety zone, or by changing its mode) to prevent the incompatibility. The zone in which the responsibility must be transfered is called the *consistency zone* of the mode m, denoted $CZ(m)$.

Critical coverage If the responsible entity chooses to transfer its responsibility to ensure that all incoming entities have an accurate view of the state of the responsible entity when entering $CZ(m)$, timely communication must be guaranteed in a zone $CC(m)$ around $CZ(m)$. This corresponds to the *critical coverage* associated with mode m of the responsible entity. Upon failure of communication (i.e., when the critical coverage of m is not covered), a responsible entity needs to adapt its behaviour, by entering a mode whose critical coverage is covered. The different zones are illustrated in Figure 2.

4.3.3 Deriving requirements on entities' behaviour

Requirements on entities' behaviour can be systematically derived from the contracts between entities. These requirements describe the conditions that an entity must fulfil before transitioning to or remaining in a mode. These conditions are expressed in terms of:

- a zone in which an entity must be able to communicate to switch to or remain in a mode (i.e., critical coverage for this mode),
- the advance warning it must give to other entities before switching to a given mode (i.e., the value of t_{warning}),
- the reaction it must have when receiving a message, depending on the content of the message (typically, switching to a mode and/or sending a message).

The conditions can all be automatically derived from the contract type, as a function of the contract parameters (e.g., t_{warning} or t_{feedback}), the values of which are application-specific.

5 Emergency vehicle warning system

In this section, we show how an emergency vehicle warning system using real-time feedback can be designed with Comhordú. The system returns the maximum speed at which the emergency vehicles can drive to ensure road safety given currently available communication. When timely communication cannot be guaranteed, the new speed at which it is safe to drive is provided to the emergency vehicles in real-time so that they can negotiate traffic safely as some vehicles might be in their way. Our emergency vehicle warning system therefore allows to reduce delays when the quality of communication is sufficient, while guazranteeing that no accident happens as long as the requirements are respected.

5.1 Specifying the safety constraints

The first step in using the coordination model is to formalise the safety constraints of the emergency vehicle warning system.

5.1.1 Scenario, modes and states

In the emergency vehicle early-warning system, an emergency vehicle informs other vehicles of its arrival. Upon reception of such a message, vehicles attempt to free the way for the emergency vehicle, and if not possible, send a message back to the emergency vehicle to inform it of the situation. The emergency vehicle can therefore regulate its speed depending on the achievable communication and, potentially, feedback from other vehicles.

As previously discussed, there are two types of entities in this system: emergency vehicles and other vehicles, and the goal is for emergency vehicles to drive as fast as possible to their destination while the goal of the other vehicles is to make progress safely towards their destination. The modes of the emergency vehicles are the following: `stopped`, $\{\texttt{going\_at\_v}_i, \texttt{accelerating\_to\_v}_i, \texttt{braking\_to\_v}_i\}_{i\in[1,n]}$; and the other vehicles can be modelled with the following modes: `travelling`, `getting_out_of_the_way`, and `out_of_the_way`. The behaviour of the two entities and the transitions between their modes are shown in Figures 3 and 4. The states of the two entites can be described by their mode, their position and their speed: `Mode x Position x Speed`.

With these definitions, the safety constraint that emergency vehicles must not crash into other vehicles at any time can be expressed as a condition on the state of the entities, as shown in Equation 1 (see Section 4.1.3).

5.2 Ensuring the safety constraints

Once the safety constraints have been formalised, the compatibility between the modes of the entities can be deduced. Table 2 presents the mode compatibility and shows that `stopped` and `out_of_the_way` are the fail-safe modes of the emergency vehicle and the other vehicles respectively.

To allow vehicles to make progress when they are not in the vicinity of emergency vehicles, emergency vehicles are made responsible for preventing incompatibilities. Two types of contract are then possible.

Table 2 Mode compatibility

Other Vehicle	Emergency Vehicle: stopped	acc. _to_v$_i$	brak. _to_v$_i$	going _to_v$_i$
travelling	✓	✗	✗	✗
getting_out _of_the_way	✓	✗	✗	✗
out_of_the_way	✓	✓	✓	✓

Contract without feedback If we choose a contract without feedback between the emergency vehicle and the other vehicles, the emergency vehicle has to warn the other vehicles early enough before its arrival, so that they will get out of its way. For this purpose, the emergency vehicle sends a transfer message containing the current time, position and velocity of the emergency vehicle, and a description of what action it is planning to take (e.g., to accelerate): `<time, position, speed, actionPlanned>`. The other vehicles use the information in the message to get out of the way. However, if communication degrades and the coverage goes below the critical coverage of its current mode, the emergency vehicle is notified within t_{feedback} that it should slow down. By reducing its speed, the emergency vehicle will enter a mode whose critical coverage is currently covered and will therefore remain able to notify other vehicles that it arrives, and this at least t_{warning} in advance. If the size of the actual coverage eventually reaches zero, then the emergency vehicle has to stop.

Contract with feedback Alternatively, a contract with feedback can be chosen for the emergency vehicle warning system. This contract enables a vehicle to reply to the transfer message from the emergency vehicle if it prefers not to get out of the way. Note that it is unlikely that this vehicle cannot communicate with the emergency vehicle, as it has just received a message from it. If it cannot reply, then it has to get out of the way. Upon reception of a feedback message from a vehicle, the emergency vehicle slows down until the vehicle gets out of the way or until the emergency vehicle finds another route to reach its destination.

In both contracts, when communication improves, the emergency vehicle is informed in real-time of its new maximum speed and can accelerate safely.

A contract with feedback is more suited for this system because it does not rely on the assumption that vehicles are always able to get out of the way of emergency vehicles.

5.3 Deriving requirements on entities' behaviour

Comhordú allows to derive the requirements on the behaviour of emergency vehicles and other vehicles automatically, given a chosen contract. For example, the size of the consistency zone CZ_i for the modes going_at_v$_i$, accelerating_to_v$_i$, and braking_to_v$_{i-1}$ can be derived:

$$CZ_i = SZ + t_{\text{warning}} \cdot v_i \quad (2)$$

with $SZ = d$, the safety distance that emergency vehicles must maintain to other vehicles. This derives from the fact that the contract between vehicles and emergency vehicles states that emergency vehicles must warn vehicles at least t_{warning} before arriving at their location.

Similarly, the expression of the critical coverage CC_i of the modes `going_at_vi`, `accelerating_to_vi`, and `braking_to_vi-1` can be derived:

$$CC_i = (present + period) \cdot v_i + \max\Big(CZ_i, \\ \big(adaptNotif + R\_reaction(v_i)\big) \cdot v_i + CC_{i-1}\Big), \quad (3)$$

where $present$, $period$, and $adaptNotif$ are the space-elastic model parameters, and $R\_reaction(v_i)$ is the time required for an emergency vehicle to brake from v_i to v_{i-1}. This illustrates that emergency vehicles travelling at speed v_i must send messages to a zone that is big enough to ensure that (i) other vehicles have time to receive a message before entering the consistency zone, (ii) if a message is not delivered, the emergency vehicle has time to be notified and slow down, before sending another message.

Comhordú also allows to derive requirements for an emergency vehicle to accelerate safely from v_i to v_{i+1}: the critical coverage CC_{i+1} must be covered, and it must warn vehicles at least:

$$\Delta = msgLatency + \max(t_{\text{warning}}, adaptNotif) \quad (4)$$

in advance, where $msgLatency$ and $adaptNotif$ are parameters of the space-elastic model. Thus, an emergency vehicle must warn vehicles when it intends to accelerate, early enough so that they have time to avoid it, and that it can cancel accelerating if the message has not been delivered in a big enough zone.

These conditions can be expressed on the transitions between modes, see Figure 5 and 6.

To summarise, the emergency vehicle warning system provides a speed v_i whose CC_i is currently covered. When the coverage increases and becomes greater that CC_{i+1}, it waits for a message sent at least Δ ago before accelerating to v_{i+1}. Upon reception of a transfer message from an emergency vehicle, vehicles must ensure that they do not collide into it, by either getting out of its way within t_{warning} or sending a feedback message within $t_{\text{feedback}} - msgLatency$. The emergency vehicle should stop within $t_{\text{warning}} - t_{\text{feedback}}$ after receiving the feedback. The values of the contract parameters t_{warning} and t_{feedback} must be carefully chosen so that the aforementioned conditions are respected.

5.4 Evaluation

We have presented the requirements on the behaviours of both emergency and other vehicles to ensure that emergency vehicles can safely rely on wireless communication to travel faster. Wireless communication, especially when it is multi-hop, i.e., relayed by other vehicles or the infrastructure, can cover a larger range. For example, experiments with an implementation of the space-elastic model (Hughes, 2006) have shown that low jitter can be guaranteed for delivery over 3 hops, which can correspond to up to 750 m using IEEE 802.11b. Using Equation 2 and 3, and the values $t_{\text{warning}} = 30\,\text{s}$, $period = 2\,\text{s}$, $present = adaptNotif = 0.5\,\text{s}$, this would allow emergency vehicles to travel at a speed of 80 km/h, provided that other vehicles can get out of their way within the 30 s warning. While this might not be achieved all the time, it is a very promising result, especially as traffic lights could also use the messages to give priority to the direction of the approaching emergency vehicle.

Moreover, we compared the maximum speed that emergency vehicles can drive using a warning system with feedback relying solely on wireless communication, with the velocity that can be achieved with a similar warning system without feedback. As emergency vehicles do not have guarantees that their messages are received in a warning system without feedback, they have to use audible or visible warnings and their speed is limited by their line of sight. Therefore, for the warning system without feedback, we distinguished urban environments (UE) where the driver's view is likely to be obstructed from rural environments (RE) where the horizon is more likely to be unobstructed. The warning system with feedback should however function similarly in urban and rural environments since it does not rely on the driver's horizon.

For this evaluation, we used 120 km/h for the maximum velocity of emergency vehicles ($v_{\max}$), v varies between 0 and $v_{\max}$ by increments of 10 km/h, and the other parameters have the values specified above. The warning system with feedback is expected to enable a higher velocity for emergency vehicles when the size of the actual coverage is large compared to similar systems without feedback. When the actual coverage is small, the warning system using sirens is expected to be more effective.

The variations of the maximum speed over the size of the actual coverage are depicted in Figure 7. The figure shows that when the coverage is above 350 m, emergency vehicles using feedback can drive faster than without feedback in urban environments. In rural environments, however, the threshold is slighty higher at 850 m. The slackening of the rate of increase of speed for emergency vehicles using feedback above 90 km/h is due to the delivery zone required to react safely to an adaptation notification becoming bigger than the consistency zone (c.f., Equation 3).

Our evaluation shows that using feedback allows emergency vehicles to travel faster when communication is sufficient. This system used in conjunction with traditional systems allows emergency vehicles to gain precious seconds. The main shortcoming of our emergency vehicle early-warning system is the assumption that vehicles are all equipped with wireless communication capabilities, but with the widespread adoption of inter-vehicle communication (Khaled et al., 2005) we envision that this limitation will soon be overcome.

6 Conclusion

In this paper, we presented a novel emergency vehicle early-warning system that enables emergency vehicles to use new communication technologies to travel faster, which can be crucial in emergency situations. The use of a real-time communication model with feedback allows emergency vehicles to rely on wireless communication when it is sufficiently good. In these conditions, they send messages to vehicles on their path to warn them of their arrival, and because the range of wireless communication can be significantly higher than the line of sight, emergency vehicles can travel faster. When the quality of real-time communication is not sufficient, emergency vehicles are notified and can slow down and revert to the traditional method of sirens and beacons, hence ensuring the safety of both emergency vehicles and other vehicles.

AKNOWLEDGEMENTS

The authors are grateful to Science Foundation Ireland for their support of the work described in this paper under Investigator award 02/IN1/I250 between 2003 and 2007.

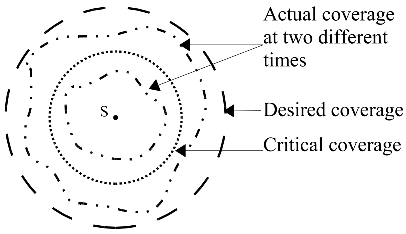

Figure 1 Different coverages of the space-elastic model

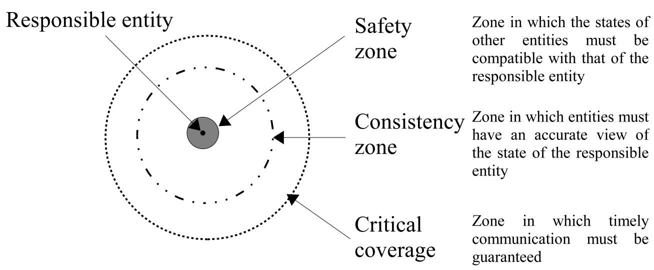

Figure 2 Definitions of the different zones

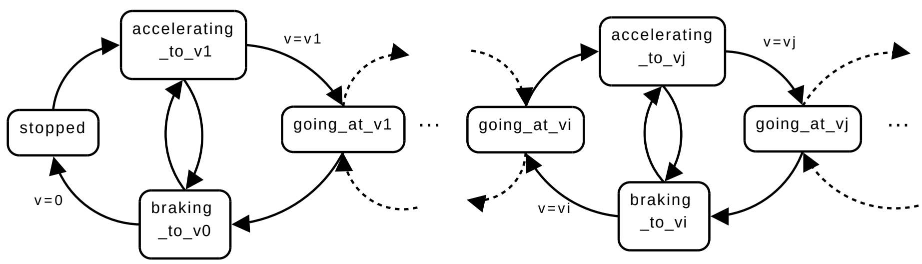

Figure 3 Modes of emergency vehicles

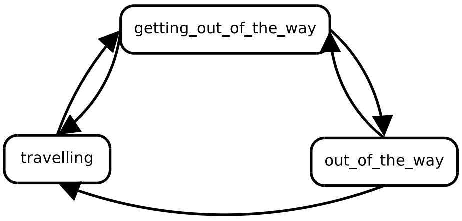

Figure 4 Modes of vehicles

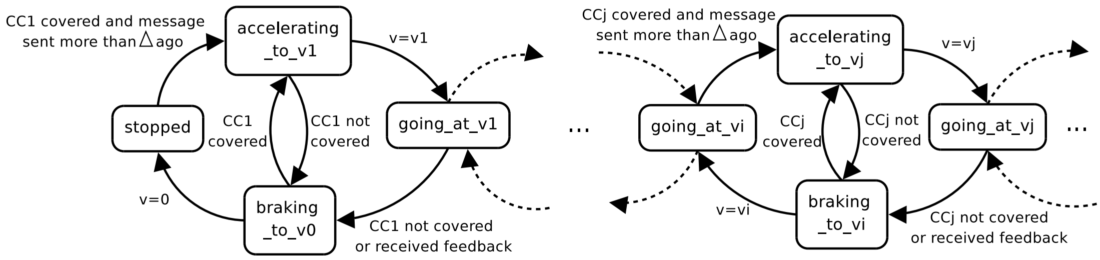

Figure 5 Modes of emergency vehicles with derived requirements

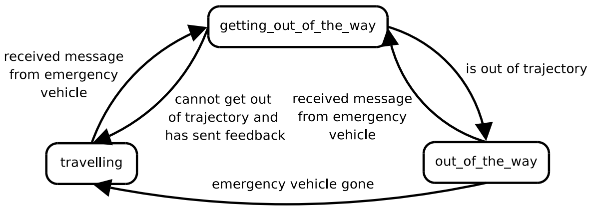

Figure 6 Modes of vehicles with derived requirements

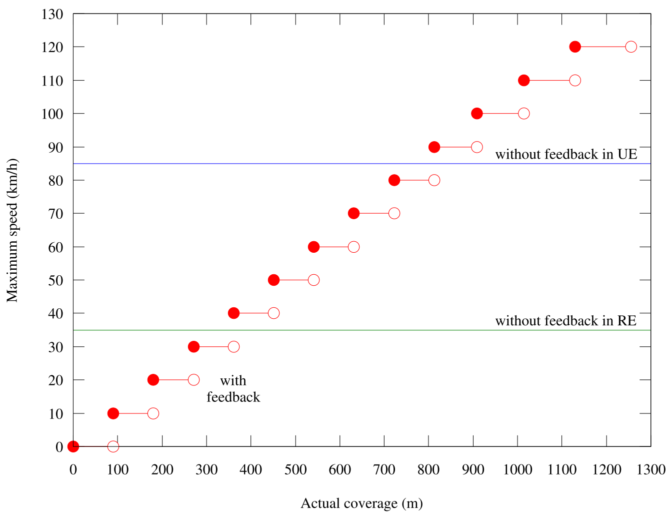

Figure 7 Achievable speed versus actual coverage size

References

Baber, J., Kolodko, J., Noel, T., Parent, M., and Vlacic, L. (2005). Cooperative autonomous driving: Intelligent vehicles sharing city roads. *IEEE Robotics & Automation Magazine*, 12(1):44–49.

Bouroche, M., Hughes, B., and Cahill, V. (2006a). Building reliable mobile applications with space-elastic adaptation. In *Mobile Distributed Computing workshop (MDC 2006)*.

Bouroche, M., Hughes, B., and Cahill, V. (2006b). Real-time coordination of autonomous vehicles. In *IEEE Conference on Intelligent Transportation Systems (ITSC 06)*.

Brown, L. H., Whitney, C. L., Hunt, R. C., Addario, M., and Hogue, T. (2000). Do warning lights and sirens reduce ambulance response times? *Prehospital Emergency Care*, 4(1):70–74.

Casturi, R. (2000). *A Macroscopic Model for Evaluating the Impact of Emergency Vehicle Signal Preemption on Traffic*. PhD thesis, Faculty of Virginia Polytechnic Institute.

Code, K., Image, R., McKenna, L. H., Images, V. P., and Class, P. (2006). Emergency warning system for approach of right of way vehicle. United States Patent.

Corporation, R. T. (2002). Emergency vehicle automatic crash notification and event reporting technology. White Paper.

Dötzer, F., Kohlmayer, F., Kosch, T., and Strassberger, M. (2005). Secure communication for intersection assistance. In *2nd International Workshop on Intelligent Transportation*, Hamburg, Germany.

Durresi, M., Durresi, A., and Barolli, L. (2005). Emergency broadcast protocol for inter-vehicle communications. In *11th International Conference on Parallel and Distributed Systems*, volume 2, pages 402–406.

Eisenberg, M. S. (1998). Defibrillation: the spark of life. Scientific American.

Ewing, J. L. and Zubiel, S. T. (2006). Emergency vehicle warning system. United States Patent.

Farkas, K., Heidemann, J., Iftode, L., Kosch, T., Strassberger, M., Laberteaux, K., Caminiti, L., Caveney, D., and Hada, H. (2006). Vehicular communication. *IEEE Pervasive Computing*, 5(4):55–62.

Hughes, B. (2006). *Hard Real-Time Communication for Mobile Ad Hoc Networks*. PhD thesis, Dept. of Computer Science, Trinity College Dublin.

Julien, C. and Roman, G.-C. (2007). Egospaces: Facilitating rapid development of context-aware mobile applications. *IEEE Transactions on Software Engineering*. To appear.

Khaled, Y., Ducourthial, B., and Shawky, M. (2005). Ieee 802.11 performances for inter-vehicle communication networks. In *IEEE Vehicular Technology Conference*, volume 5, pages 2925–2929.

Kolodko, J. and Vlacic, L. (2003). Cooperative autonomous driving at the intelligent control systems laboratory. *IEEE Intelligent Systems*, 18(4):8–11.

Lee, W.-H., Jeng, B.-S., Tseng, S.-S., and Wang, C.-H. (2004). Electronic toll collection based on vehicle-positioning system techniques. In *IEEE International Conference on Networking, Sensing and Control*, volume 1, pages 643–648.

Little, T. D. C. and Agarwal, A. (2005). An information propagation scheme for vehicular networks. In *IEEE Intelligent Transportation Systems Conference*, Vienna, Austria.

Mayhew, G. L. and Shirkey, K. L. (1992). Communication architectures supporting proactive driver safety. In *Intelligent Vehicles Symposium*, pages 370–377, Detroit, MI, USA.

Miyawaki, M., Yamashiro, Z., and Yoshida, T. (1999). Fast emergency pre-emption systems (fast). In *International IEEE Conference on Intelligent Transportation Systems*, pages 993 – 997.

Mock, M. (2004). *On the Real-Time Cooperation of Autonomous Systems*. Fraunhofer Series in Information and Communicatoin Technology. Shaker Verlag, Aachen.

Murphy, A. L., Picco, G. P., and Roman, G.-C. (2001). Lime: A middleware for physical and logical mobility. In *Proceedings of the International Conference on Distributed Com puting Systems (ICDCS'01)*, pages 524–533. IEEE Computer Society.

NASA (2001). Nasa technology tested on california roads. NASA Tech Briefs.

Nett, E., Gergeleit, M., and Mock, M. (2001). Mechanisms for a reliable cooperation of vehicles. In *Proceedings of the IEEE International Symposium on High-Assurance Systems Engineering (HASE'01)*, pages 75–81. IEEE Computer Society.

Pallottino, L., Scordio, V. G., Frazzoli, E., and Bicchi, A. (2007). Decentralized cooperative policy for conflict resolution in multi-vehicle systems. *IEEE Transactions on Robotics*. To appear.

Santos, R. A., Edwards, R. M., and Seed, N. L. (2004). Supporting inter-vehicular and vehicle-roadside communications over a cluster-based wireless ad-hoc routing algorithm. In *International Symposium on Information and Communication Technologies*, pages 1–6, Cancun, Mexico.

Schemmer, S., Nett, E., and Mock, M. (2001). Reliable real-time cooperation of mobile autonomous systems. In *Proceedings of the Symposium on Reliable Distributed Systems (SRDS 2001)*, pages 238–246. IEEE Computer Society.

Transportation Research Board (2006). Traffic signal preemption for emergency vehicles: A cross-cutting study. Transportation Research Board of the National Academies.

Tsugawa, S. (2002). Inter-vehicle communications and their applications to intelligent vehicles: an overview. In *IEEE Intelligent Vehicle Symposium*, volume 2, pages 564– 569.

Xu, Q., Mak, T., Ko, J., and Sengupta, R. (2004). Vehicle-to-vehicle safety messaging in dsrc. In *1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 19–28, Philadelphia, PA, USA.

Yin, J., ElBatt, T., Yeung, G., Ryu, B., Habermas, S., Krishnan, H., and Talty, T. (2004). Performance evaluation of safety applications over dsrc vehicular ad hoc networks. In *First ACM Workshop on Vehicular Ad Hoc Networks*, pages 1–9, Philadelphia, PA, USA.