Power-adaptive intelligent broadcast using vehicular ad hoc networks for better traffic safety applications

Bür, Kaan; Kihl, Maria

2013

Document Version:
Early version, also known as pre-print

Link to publication

Citation for published version (APA):
Bür, K., & Kihl, M. (2013). Power-adaptive intelligent broadcast using vehicular ad hoc networks for better traffic safety applications. Unpublished.

Total number of authors:
2

Power-Adaptive Intelligent Broadcast using Vehicular Ad Hoc Networks for Better Traffic Safety Applications

Kaan Bür[1], Maria Kihl

Department of Electrical and Information Technology
Lund University
P.O. Box 118, 221 00 Lund, Sweden
{kaan.bur, maria.kihl}@eit.lth.se

Abstract

Vehicular communications is a challenging application area for mobile ad hoc networks. Vehicle collision warning systems are a particularly important application of vehicular communications, since traffic accidents cause hundreds of thousands of fatalities and injuries world-wide every year. One popular method for collision warning systems is based on selective broadcasting. Its popularity stems from its ability to deliver a warning message to all vehicles within a certain zone of interest without flooding the network with too many packets, thanks to its 'selective' decisions regarding packet relaying. However, implementation of vehicle collision warning systems requires close to 100% success rate, even under extremely unfavourable conditions. Such conditions arise in sparse networks, for instance, where node connectivity is low, and message dissemination becomes very difficult. Additional measures therefore need to be developed and implemented in order to keep all nodes informed. Here we demonstrate how the collision warning delivery ratio in sparse networks can be improved, using a power-adaptive intelligent broadcast algorithm we developed for information dissemination in vehicular ad hoc networks. The improvement obtained from our algorithm is based on a gradual increase of a vehicle's transmission power with each repeated message. To evaluate the effectiveness, efficiency, convergence time and power consumption of our algorithm, we compared it to two other, non-adaptive, selective broadcast protocols, and to flooding. The analysis revealed that our adaptive protocol yields a significantly higher warning delivery success rate, particularly in sparse networks, where the other algorithms suffer from the lack of connectivity. Our findings provide insight into the design requirements of high-performance intelligent broadcast algorithms.

Keywords:
Information dissemination, intelligent transportation systems, traffic safety applications, vehicular communications, wireless transmission power.

1. Introduction

THE *zero-accident* vehicle, a vehicle that cannot be involved in an accident, is both a vision and a challenge for all automotive actors worldwide. To deserve their name, zero-accident vehicles require active traffic safety systems that help the driver in hazardous situations, well in advance of the actual collision. Annual traffic statistics indicate the potential benefit of intelligent driver support systems. Statistics show, for instance, that rear-end collisions are the most common type of all accidents. Furthermore, 90% of rear-end collisions occur on straight roads; in nearly 70% of the cases, the lead vehicle is stationary at least for 1 s before the collision; and more than 90% of the drivers report they have been distracted just before the collision [1]. The automobile industry is, therefore, on the verge of a paradigm shift in vehicle design, with a need to equip cars with environmental sensing, machine learning, and mobile communication capabilities, using the latest technology in electronics, computer science, and telecommunications.

The primary objective of traffic safety applications is to save lives by preventing accidents, while also reducing congestion, cost of operation, and air pollution. Being potentially life-critical, they have stringent delay and reliability requirements [2], [3]. Due partly to their performance needs, and partly to their unpredictability in terms of time, space, and severity, safety applications can arguably be served better by flexible, impromptu communication technologies like vehicular ad hoc networks (VANET). In this regard, IEEE 802.11p has been standardised as a wireless access technology utilised by IEEE 1609, the system architecture for Wireless Access in Vehicular Environments (WAVE) [4]. The main advantages of IEEE 802.11p are that its equipment is cheap and standardised, and that the ad hoc communication mode may provide the low latency required for safety-critical applications. However, early customers will benefit only little from the applications, since ad hoc communications requires a certain penetration before becoming fully operational.

Manuscript submitted March 8, 2013, to *IEEE Access*.

This work is partly funded by Excellence Center at Linköping-Lund in Information Technology (eLLIIT). Maria Kihl and Kaan Bür are members of Lund Center for Control of Complex Engineering Systems (LCCC), a Linnaeus Center at Lund University, funded by the Swedish Research Council (Vetenskaprådet).

[1] Corresponding author. *Tel:* +46-46-222 7341. *Fax:* +46-46-129 948.

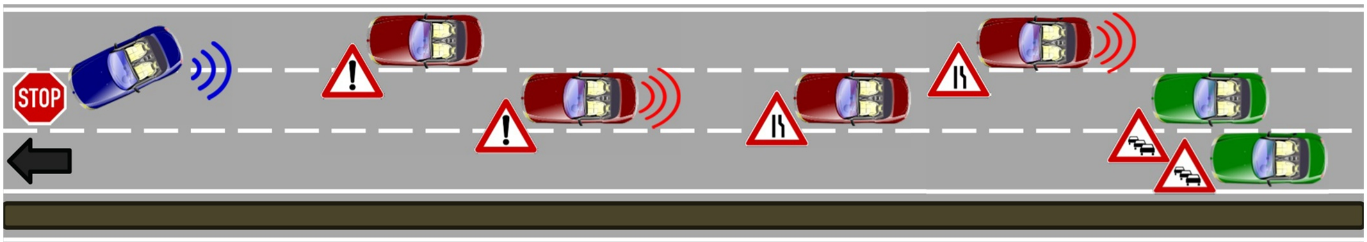

Figure 1. Active traffic safety: An emergency application utilises vehicular communications to avoid a typical rear-end collision situation.

Due to both low initial market penetration as well as the big variety in car traffic, sparse networks are not rare in vehicular communications. Lack of connectivity due to low node density prevents applications from achieving a high message delivery ratio. This is an important problem for the performance of traffic safety applications, and many existing solutions fail to address it satisfactorily. In this study, we hence focus on active safety applications in sparse VANET, and propose an algorithm adapting the cars' wireless transmission power to node density, and then increasing the power incrementally if warnings do still not get delivered. Our power-adaptive approach improves the warning success ratio significantly in sparse VANET, whereas it also works well with dense networks as it decreases the transmission power as node density becomes higher.

In this article, we first introduce the general concept of selective broadcast for safety message dissemination. Next, in light of related research, we identify some of the problems with the implementation of the concept, and present a solution proposal, comparing its performance to a selection of existing solutions. Finally, we conclude our article with remarks on future research directions.

2. Selective Broadcast in VANET

The mission-critical nature of safety applications necessitates an efficient message dissemination method which, upon the initiation of an emergency, can address all the vehicles in a given region successfully, while not generating too much communication overhead, blocking all other, potentially important, applications. Figure 1 depicts the active traffic safety concept, utilising vehicular communications in a rear-end collision case on a straight highway segment. The leading vehicle broadcasts the initial warning, which is propagated backwards by the other vehicles. Depending on the distance to the emergency site as well as individual car positions and speeds, the type of warning may vary.

The simplest form of message dissemination in this respect is flooding [5], where the emergency initiator starts a periodic warning message sequence, and the repeaters just apply the following rule:

(1) Start with your own periodical warning sequence upon receiving your first warning.

Although quite successful in propagating a message, flooding generates too much communication overhead. It has no rules to limit the number of repeaters, and no termination condition for the sending of the periodic warnings. Therefore, selective broadcast mechanisms have been devised [6], [7], enabling the receivers of the emergency message to decide whether it is necessary for them to repeat the message, thus limiting the overall message overhead. Assuming an emergency zone with cars moving in a single direction, the rule above can be extended to help the decision on message relaying:

(2) Repeat only if the message comes from the front.

This new rule ensures that the warning propagates against the direction of the traffic flow, starting at the emergency site and moving towards the vehicles approaching it, which limits the number of repeaters to one direction only. Nevertheless, it can do better if we extend it further as follows [8]:

(3) Wait before you start repeating, and cancel if you receive the message a second time while waiting.

Now, potential repeaters wait for a short time before starting their warning sequence. If, while waiting, they overhear another vehicle repeat the same warning, they cancel their sequence since another car has taken care of it. There is no need to repeat the warning any more.

The waiting time t_{wait} can simply be set randomly, where $t_{min} \leq t_{wait} \leq t_{max}$, and t_{min} and t_{max} are the minimum and maximum waiting times, respectively. Alternatively, it can be set inversely proportional to the distance between the potential repeater and the vehicle it has received the warning from, favouring as repeaters those vehicles farthest away from the sender, increasing thus the one-hop progress the warning makes. There is even a subcategory of selective broadcast algorithms which favours those nodes nearer to the edge of the sender's transmission range by adding a counter-based probabilistic element to the decision process [9].

The next idea is to reapply the third rule to those nodes already actively sending periodic messages:

(4) Stop your on-going periodic warning sequence if you receive your own message from your back.

This last rule provides the warning initiator and the active repeaters with a stopping condition. When, in multi-hop message dissemination, a node overhears its own message relayed by another node, it knows the message has propagated successfully. The phenomenon

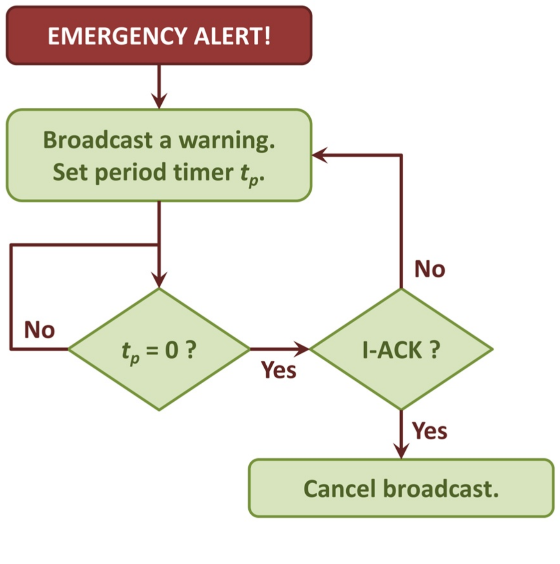

(a) Emergency initiator, starting the process

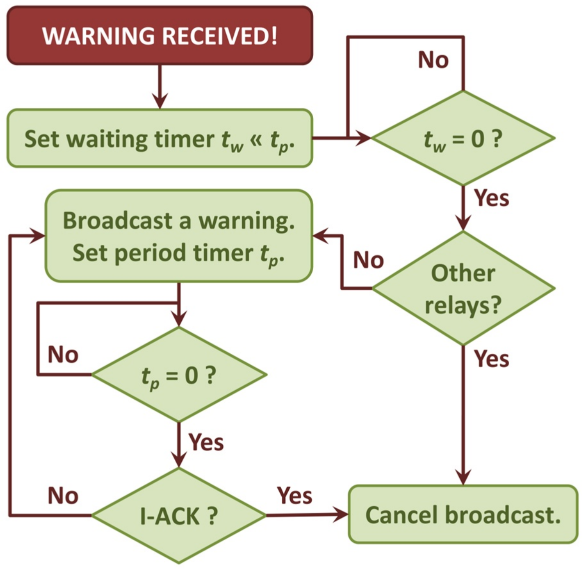

(b) Warning receiver, handling the relaying

Figure 2. Selective broadcast algorithm run for message dissemination by: (a) the initiator node; (b) the repeater nodes.

is called implicit acknowledgement (I-ACK). Making use of the I-ACK mechanism is considered a means to improve reliability since, whenever a sender fails to receive an I-ACK, it knows that retransmission in the next period is necessary.

Putting these rules together [6], [8], the first vehicle encountering an emergency situation initiates a periodic warning sequence. It stops sending messages as soon as it overhears another vehicle at its back propagating the warning. Upon receiving a message coming from their front, other vehicles start their own periodic broadcast sequence, provided that they do not overhear another vehicle starting before they do. The repeaters, too, stop once they overhear the warning being propagated. Figure 2 shows how selective broadcast is operated at (a) the warning initiator; and (b) the repeater nodes.

In the following section, we discuss the shortcomings of selective broadcast. We find the discussion important since it provides us with a set of design considerations for a high-performance selective broadcast algorithm.

3. Discussion on Algorithm Improvements

Our previous evaluation of selective broadcast [10] has shown us that it performs generally well in terms of information dissemination. However, there are some problems it cannot cope with, the most important being low network connectivity. In VANET, as in mobile ad hoc networks, coverage is a function of the number of a node's neighbours. It is therefore not surprising that success rates are low under low network density.

A special situation in which the gaps in connectivity affect the performance even worse is when the car keeping the warning alive runs away. The algorithms usually assume that all cars stop in an emergency. In real life, this is not always the case, and scenarios are possible where "not all the cars stop". In such a case, depicted in Figure 3, a relay node starts repeating the warning. The initiator, having received I-ACK, cancels its own. The repeater then moves away. The cars at the back, being outside the repeater's transmission range, never get the warning. We call this phenomenon *the runaway repeater syndrome*.

One possible solution is to reactivate one of the former repeaters' relaying function as soon as an active repeater overtakes it. The question here is how to make sure that our selection solves the problem. Even if we set a rule for selecting a node to reactivate (e.g. pick the last active repeater), we still don't know whether it is slow enough. We need at least one node to remain on the emergency site in order to keep the message alive. An obvious candidate is the original emergency initiator. Having stopped broadcasting upon overhearing its own message being relayed, it needs to be reactivated as soon as it is overtaken by an active repeater.

The reactivation of the emergency initiator actually turns the *runaway repeater* problem into the *runaway emergency initiator* problem. There is, however, another rather obvious solution: adaptive transmission power. Even a runaway repeater can reach a receiver if it increases its transmission power once it realises that

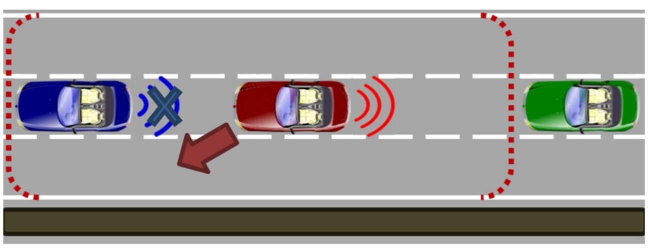

Figure 3. Runaway repeater syndrome: The relay node is too far from the others; the red car runs away with the warning.

there are no cars at its back keeping the warning alive.

Concerning the repeater selection, it seems sufficient to have a criterion to keep the number of repeaters under control; it does not make a crucial performance difference to select the best one [10]. In this regard, introducing a waiting time before relaying a message works quite well, and the distance-based delay is not particularly better than the random delay. This is most likely because the information on the distance between the sender and the receiver, which the waiting time is based on, becomes obsolete just after the first use.

Finally, periodicity improves the warning success rate significantly by creating data redundancy; while the stopping condition effectively prevents flooding [10]. Counter-based protocols tend to suffer from broadcast probability calculation. Much time needs to be spent until a sufficient number of messages is accumulated to make the probability calculation work well.

For these reasons, we take the selective broadcast algorithms which are not counter-based but periodic, and adopt the much simpler random waiting time approach, as our starting point for our improvements. We address the fundamental issue of low network connectivity in the next section, proposing a power-adaptive solution to the problem.

4. Power-Adaptive Intelligent Broadcast

The importance of adapting the transmission power to network conditions has repeatedly been pointed out in the research literature [6], [11], [12], [13], [14], as a general method of more efficient communication. We can think of two important, and usually conflicting, reasons for power adaptation: congestion control and network connectivity [11]. Good congestion control requires that the transmission range is kept low to minimise contention, interference, and packet collision, thus avoiding channel saturation, especially in dense networks [6]. This also increases the reliability of the overall system in terms of packet reception rate [11], [12]. When the network is sparse, on the other hand, network connectivity improves if the transmission range is increased [13]. Here, node connectivity is the predominant factor determining network performance, rendering congestion and collision relatively rare events. Thus, the optimisation of the transmission range is about finding the balance between good coverage and congestion control [11].

Typically, the adaptation of the transmission power can be based on the node density [12], [13], defined as:

$$\frac{n_{Nbr}}{N} = \frac{2 \times r_{Tx}}{l_{Road}}$$

Where n_{Nbr} is the number of a node's neighbours, N is the total number of nodes, r_{Tx} is the transmission range, and l_{Road} is the size of the relevant road segment. Some of this information is available from the beacons a node receives from its neighbours, while others are not. It is also possible, however, to define a relation between a car's own speed and the network density [12], [13], hence deriving a rule for the transmission range.

As part of our study, we want to investigate the performance of a power-adaptive intelligent broadcast algorithm in comparison with none-adaptive, selective broadcast protocols, and flooding. Our objective is to improve the warning success rate of safety applications in sparse vehicular networks. One could set all safety messaging to be transmitted at maximum power, of course, but this would increase interference and packet collision to inacceptable levels, as mentioned before [6]. Given the importance of the safety messages, it could be argued that it would still be beneficial to transmit these at maximum power, in the hope that the warning would propagate successfully even at a low packet delivery ratio. But we need to consider power efficiency as it is an important issue both for protecting the environment as well as for the new and emerging electric cars.

Adapting the transmission power solely to network density means that wireless channel dynamics, e.g. noise, contention, and random errors on the radio interface, are not considered properly and the algorithm acts purely statistically. We have thus adopted a slightly different approach, based not only on network density but also on additional information that can be extracted from the network. We start with the following rule:

(1) Estimate network density; set initial emergency transmit power accordingly.

One important source of information for our purpose is the beacons exchanged between the vehicles. We have already mentioned that beacons can help us to estimate the network density, but their primary task is to inform each vehicle on its neighbours' position, speed, and driving direction. Using this information, it is possible for a vehicle in emergency to estimate the number of neighbours within a given range, driving in the same direction and coming from the back. Hence, our approach in emergency is to set the initial transmission power according to node density, while also making small adjustments in order to ensure that a predefined minimum number of cars at the back will be reached:

(2) Estimate number of reachable cars at your back; fine-tune initial power if car count is less than a predefined threshold.

Beacons contain useful information, but they are single-hop; the sender does not know whether they have been received. In contrast, emergency warnings are multi-hop; they provide their sender with implicit acknowledgements (I-ACK) as soon as they are repeated by relay nodes. Moreover, the distance information provided by the beacons may vary and, thus, not be very accurate. Even when all nodes transmit their beacons at the same default power, link symmetry cannot be guaranteed. As a result, vehicles may have a neighbour map not exactly matching the reality. These inaccuracies can be partly eliminated by the I-ACK mechanism:

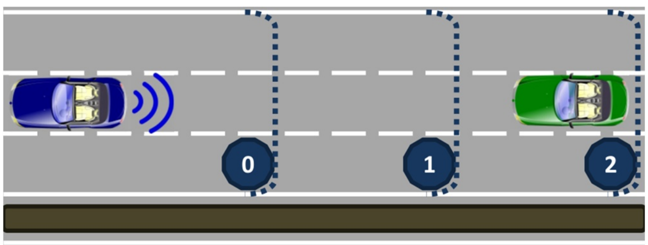

Figure 4. Power-adaptive intelligent broadcast: The leading car has a default transmission range (0), but transmits at higher power in emergency (1), and increments the power level until a car at the back are reached (2) and the warning is forwarded.

(3) *Send emergency message; wait for it to be relayed by another car; increase power if you do not hear anything before resending.*

This rule lets the sender and relay nodes increase their transmission power every time they do not overhear their own message relayed by another car at their back. The cars have multiple transmission power levels they can use, as shown in Figure 4, with the lowest predefined as the default level for their regular communication. In an emergency, the nodes start with a transmission power determined by the network density and fine-tuned by the number of cars reachable with this power. At the end of every transmission interval, as they are about to decide whether or not to retransmit their message, the cars increment the transmission power level by one if I-ACK has still not been received.

The non-reception of I-ACK is a sign showing that, with high probability, there are no other cars at the back of the sender within the current transmission range to receive the warning. Under these circumstances, increasing the transmission power should yield better packet reception rates. Unfortunately, it also causes transmission range asymmetry between one repeater and the next: Node A transmits at increased emergency power. Node B receives the warning, but relays it at the initial power. Node A cannot hear the relayed message, thus missing I-ACK, and continues sending messages, unnecessarily. We address the issue with this rule:

(4) *Reset power to initial value when you have sent a certain number of messages at maximum power.*

The rule does not really solve the transmission range asymmetry problem; it only limits the extra power consumed by the node sending messages unnecessarily. An alternate rule, which might be a more effective solution, is to limit the number of repetitions at a relay node once the maximum transmission power level is reached, and cancel message repetition completely after that. In fact, doing so will also limit excessive message traffic when a node has already tried to reach others many times at maximum power. In any case, it is worth mentioning that, even under transmission range asymmetry, the power consumption should not be too high since the algorithm keeps the number of active senders low. We discuss this and various other aspects of system performance further in the next section.

5. Performance Evaluation

In order to evaluate the performance of power-adaptive intelligent broadcast, we have compared it to two different flavours of selective broadcasting, one periodic and one counter-based, as well as to simple flooding; the details of these algorithms are described in the preceding sections. We have implemented them in ns-3 [15], release 14.1, and compared their performance through simulation. To create a realistic simulation environment, we have worked on the mobility model, physical (PHY) and medium access control (MAC) layers, and a variety of scenario parameters, such as message size and rate, number of cars and how these are initially distributed, number of lanes, node and lane speeds, road dimensions, average inter-vehicle distance, and direction of traffic. In the following sections, we first describe the emergency scenario we have simulated. Then, we explain how we have configured and customised ns-3. Finally, we introduce our performance criteria, and present the results we obtained.

5.1. Emergency Scenario

In our tests, we simulate a highway scenario with 3 lanes in a single direction, each lane having individual speed limits ranging from 60 km/h to 120 km/h, as shown in Figure 5. Vehicles are first assigned individual constant speeds, and then placed in the highway segment on the lane corresponding to their speed. In terms of *(x, y)* coordinates, the lanes represent the *y* coordinate, whereas the *x* coordinate of each vehicle is assigned randomly with uniform distribution. Two exceptions to this initial placement are the emergency initiator and the last car, placed at the beginning and the end of the segment, respectively. Once the simulation starts, the cars move along their lane according to their speed, changing their *x* coordinates. They do not change lanes, and we do not enforce safety distances since we want to create a scenario in which there are cars endangered by the distance to the cars in front. Table 1 summarises our simulation parameters and their values.

In all our scenario variants, the emergency initiator stops as soon as an emergency occurs. As for how the repeaters react upon receiving the warning, we imagine

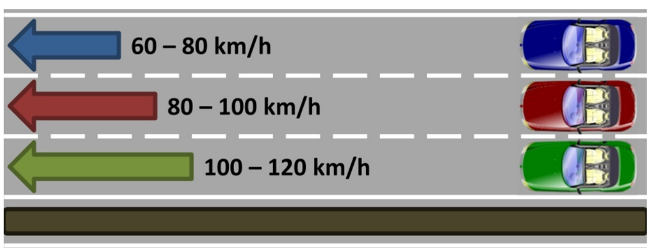

Figure 5. Simulation scenario: Lanes with different speed limits.

Table 1. Simulation parameters and settings.

Description:	Value:
Number of vehicles *(variable)*	10 .. 50
Emergency warning size	400 B
Emergency warning interval	100 ms
Wait-before-send time *(t_{min} .. t_{max})*	0 ms .. 25 ms
Beacon message size	400 B
Beacon message interval	100 ms
Beacon transmission power	100 mW
Highway segment length	1 000 m
Simulated emergency duration	10 s

Table 2. Power levels and corresponding transmission ranges.

Power Level	dBm Level	Power	Transmission Range
0 (default)	20 dBm	100 mW	≈ 160 m
1	23 dBm	200 mW	≈ 200 m
2	26 dBm	398 mW	≈ 250 m
3	29 dBm	794 mW	≈ 320 m
4 (maximum)	32 dBm	1 585 mW	≈ 400 m

3 possible variants: (1) they, too, stop; (2) they don't stop; (3) some of them, e.g. only those on the emergency lane, stop. When they stop, they do it according to their current speed, a constant driver reaction time of 1.6 s, and a constant deceleration rate of 4 m/s^2.

In terms of data traffic, we generate high-priority emergency packets and low-priority background traffic. The background traffic is represented by an application generating periodic beacons, i.e. single-hop broadcast HELLO packets, which, depending on the simulation settings, can be stopped as soon as an emergency is initiated, or can continue throughout the simulation.

5.2. Network Simulator

As our tool for discrete event simulation, we have chosen ns-3 [15], a free software licensed under GNU GPLv2 to replace eventually the very popular ns-2 [16]. ns-3 has already most of the models and functions of its predecessor and, like ns-2, is highly trusted among the network research community. Table 2 shows the 4 wireless transmission power levels we defined for the adaptive algorithm to choose from. These power levels comply with the WAVE standards in Europe and USA. More specifically, the default power level is currently set to 20 dBm, whereas a maximum of 33 dBm is defined for safety applications in both standards. As ns-3 lets us set the dBm levels only with equal integer intervals, however, we use 32 dBm as maximum power instead.

Within the node structure of ns-3, the broadcast algorithms we implement run at the network (NET) layer. Concerning the models and their parameters at the lower layers, i.e. MAC and PHY, we use Yet Another Network Simulator (YANS), also provided by ns-3. More specifically, we work with WiFi 802.11 at 6 Mbps, using OFDM at 5 GHz with 10 MHz channels. This setting enables us to model WiFi 802.11p CCH, which, as it is deprecated in ns-3, cannot be used directly. We utilise the constant speed propagation delay and log-distance propagation loss models. As required by 802.11p, we use a QoS-aware MAC model with 4 classes, although we only employ the default, best-effort class.

We have made a small change in the ns-3 YANS code for implementing the adaptive transmission power with different levels. In the original code, a lower MAC module resets the transmission power regardless of the actual input value coming from the above layers. As a workaround, we put the power level information into the broadcast packet header at NET, so the value can be read from the packet at PHY just before transmission, and the transmission power set accordingly.

5.3. Results at Default Transmit Power

Typical performance metrics used in the evaluation of communications systems, such as packet loss, error, or delivery ratios, delay, and overhead, have a pure network performance point of view. In our study, we are more interested in the performance of the system as a mission-critical safety application. Our understanding of performance is, thus, slightly different. In this section, we compare our algorithm to its competitors in terms of the following: (a) The warning delivery success ratio, i.e. the percentage of vehicles successfully receiving a warning; (b) The total number of warning messages sent per vehicle, i.e. the messaging overhead; (c) The warning convergence time, i.e. the time it takes the last vehicle in the whole group to receive a warning. The results we present are the average of 100 runs for each data point, with every run simulating a time interval of 10 s. The chosen simulation duration gives all the cars in the simulation enough time to enter a so-called *emergency zone*, in which they need to receive a warning to avoid a crash. We focus specifically on the algorithms' performance at the default power level in this section; we present our findings on performance at the maximum power level in the next section.

Figure 6 shows our simulation results regarding the effectiveness and efficiency of the algorithms. (a) The power-adaptive algorithm achieves a warning delivery success rate of 100% in even very sparse networks. The other algorithms suffer from low network connectivity and fail to deliver the warning to a significant portion of the vehicles. These algorithms recover from the gaps in connectivity and start performing better only after the network density increases to a sufficient level. (b) The power-adaptive algorithm manages to keep its message overhead significantly lower than its competitors. The other algorithms, except flooding, have a low overhead, too, but given the difference in the warning delivery ratios, it is obvious that the power-adaptive algorithm exploits the warning messages much more efficiently. (c) The power-adaptive algorithm is able to reach all the vehicles in the simulation within less than 1 s. Periodic selective broadcasting and flooding, on the other hand, cannot reach all cars when the network is really sparse, and still have a longer convergence time for relatively

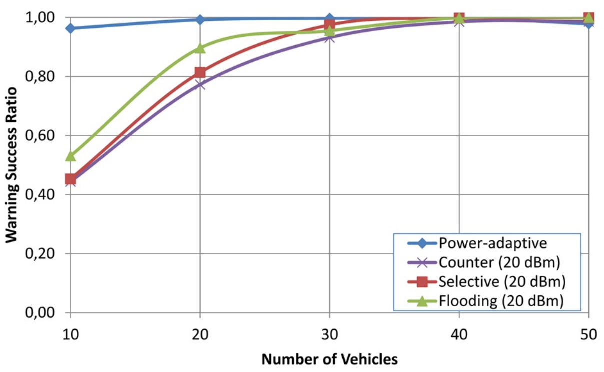

(a)

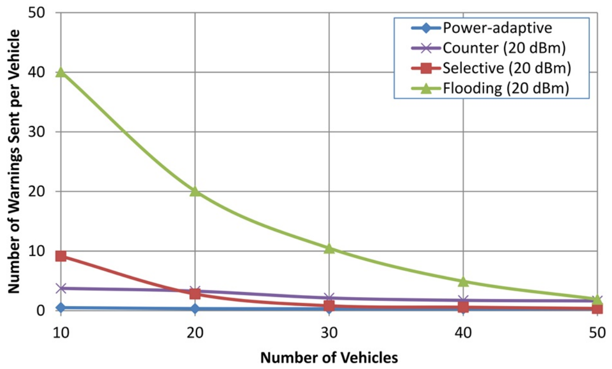

(b)

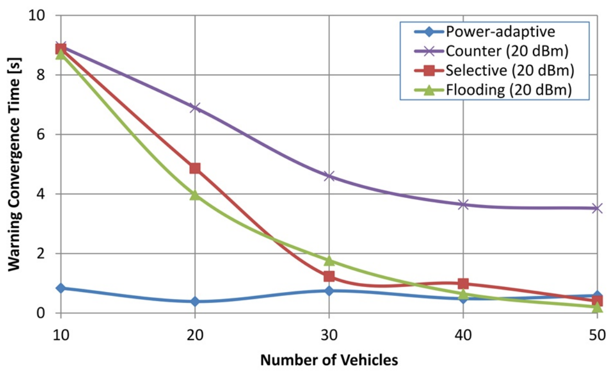

(c)

Figure 6. Simulation results comparing the effectiveness and efficiency of the algorithms: (a) Warning success ratio, i.e. ratio of vehicles successfully receiving a warning; (b) Total number of warnings sent per vehicle, i.e. message overhead; (c) Warning convergence time, i.e. interval between initiation of emergency and the time last reachable vehicle is warned.

denser networks. In case of counter-based selective broadcast, by its average convergence time we can tell that there are nearly always a few vehicles not reached throughout the simulation.

In addition to the warning success ratios we present here, we have also looked into what we call the ratio of the vehicles *warned in time*, comparing a vehicle's distance to the emergency zone, d_{emer}, as it receives the warning, to its stopping distance, d_{stop}, defined as:

$$d_{stop} = v_{init}\left(t_{react} + \frac{v_{init}}{2 \times a_{dec}}\right)$$

Where v_{init} is the speed, t_{react} is the reaction time, and a_{dec} is the deceleration rate. It should be noted that, as we define d_{emer}, we take into account the effect of the oncoming traffic, i.e. of the cars piling up behind the emergency initiator. Each new car approaching the site extends the emergency zone. Hence, d_{emer} is not merely the distance between a car and the emergency initiator; it defines the size of a region that shrinks proportionally to the number of cars having piled up previously. With this comparison, we can tell whether the warning arrived "too late" or "in time" for a particular car. The warning success ratio then becomes the percentage of the vehicles warned "in time". The results we have obtained with this criterion are actually similar to those presented in Figure 6(a), with all algorithms achieving approximately 90% of their respective performance. Although our algorithm does not have a direct influence on either d_{emer} or d_{stop}, this extra bit of information can be used by the emergency warning system to instruct the driver to brake "normal" or "hard".

5.4. Results at Maximum Transmit Power

The results presented in the previous section are all obtained at default transmission power, corresponding to the lowest level in Table 2. In other words, the competing algorithms have been limited by the default transmit power, while our power-adaptive algorithm has been free to increase its power level if necessary. As we rerun the simulations with the competitors' transmit power set to the maximum, their performance, in terms of warning success ratio, number of warnings sent per vehicle, and warning convergence time, becomes nearly equal to that of the power-adaptive algorithm. However, the competitors achieve this improvement at the cost of two important factors, which we discuss in this section.

Figure 7 compares the warning drop ratio of the power-adaptive algorithm to its competitors at (a) default; (b) maximum transmit power. The results of the adaptive algorithm are identical since its behaviour is the same in both cases. The two parts of the figure together show us the impact of transmit power on the packet loss rate. At default power, the adaptive algorithm benefits from decreasing the transmit power as the network density increases, thus achieving a lower loss rate than the regular algorithms operating at a fixed power. At maximum power, the difference in the loss rates becomes even more dramatic.

Figure 8 shows, in logarithmic scale, the algorithms' power consumption at (a) default; (b) maximum transmit power. Here, too, the results of the power-adaptive algorithm are identical since its behaviour is the same. It starts off with a high transmit power for very sparse networks, which enables it to reach all the

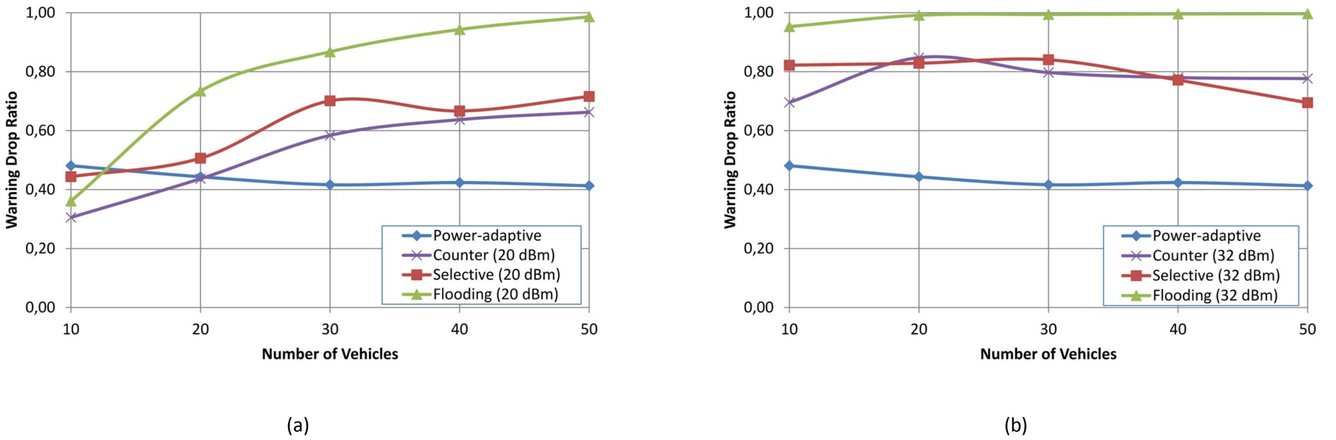

Figure 7. Simulation results comparing the warning drop ratio of the algorithms: (a) as the competitors' transmit power is fixed at the default value; (b) as the competitors' transmit power is fixed at the maximum value.

cars with a relatively low number of warnings, keeping its power consumption lower than its competitors. As the network becomes denser, the adaptive algorithm lowers its power level, but still reaches all vehicles. By adopting this strategy, it is able to find a suitable power level for all network densities to deliver the warning to all vehicles with acceptable overhead. The other algorithms have significantly higher power consumption per warned vehicle than the power-adaptive algorithm. Similar results are obtained as we look at the power consumption per sending vehicle.

The simulation results presented in this section show that, as the transmission power is increased to a level where connectivity ceases to be a problem, the other algorithms manage to improve their warning success ratio, but they pay a high price in terms of warning drop ratio and power consumption, when compared to their competitor with the adjustable transmission power strategy.

6. Conclusion

Traffic safety applications utilising VANET are one of the most challenging areas in vehicular communications research, not only due to the technological difficulties but, maybe more importantly, because of the life-critical nature of the systems under consideration. For instance, the cars participating in a traffic safety application must not fail to deliver messages as a result of the gaps in network connectivity, be it just due to time of day or due to low initial market penetration. Network sparseness, together with low wireless transmission range, prevents nodes from performing critical tasks.

In this article, we have given an overview of selective broadcasting for message dissemination in VANET and discussed its performance in terms of reliable warning delivery. We have also introduced a power-adaptive intelligent broadcast algorithm to overcome the essential problem of low network density common to the existing message dissemination schemes. The vehicles in our algorithm start at an initial transmission

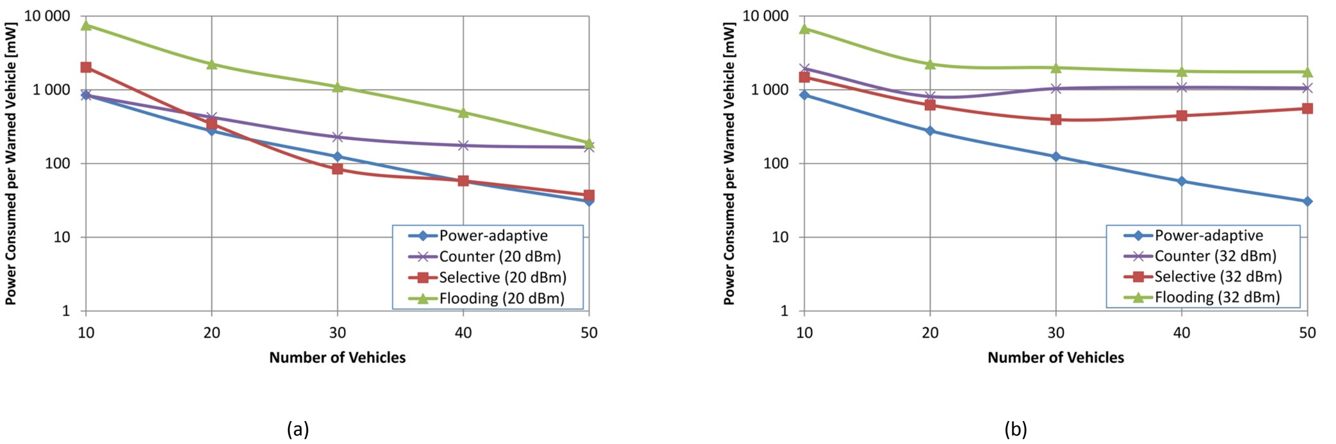

Figure 8. Simulation results comparing, in logarithmic scale, the power consumption of the algorithms: (a) as the competitors' transmit power is fixed at the default value; (b) as the competitors' transmit power is fixed at the maximum value.

power, adapted to the network density and neighbour distance, and increase their power step by step at each transmission interval, based on the non-reception of their messages repeated by others. We have compared power-adaptive intelligent broadcast to two other, non-adaptive, selective broadcast protocols, and to flooding. We have shown that the power-adaptive method performs significantly better in sparse networks.

Much work needs to be done before we can claim without hesitation that a vehicular safety warning system, based on wireless ad hoc communications, can deliver 100% reliable performance. Nevertheless, the results we achieved are promising, and help us to analyse better the design requirements of a high-performance intelligent broadcast algorithm.

Acknowledgements

We would like to thank Katrin Sjöberg, serving ETSI TC ITS as an expert for the standardization of vehicular communications in Europe, for our fruitful discussions on the ITS standards for vehicular safety applications.

References

[1] M. Brännström, J. Sjöberg, E. Coelingh, ”A situation and threat assessment algorithm for a rear-end collision avoidance system”, *Proceedings of IEEE Intelligent Vehicles Symposium*, pp. 102-107, Eindhoven, Holland, 4-6 June 2008.

[2] M. Sepulcre, J. Gozalvez, "On the importance of application requirements in cooperative vehicular communications", *Proceedings of International Conference on Wireless On-Demand Network Systems and Services*, pp. 124-131, Bardonecchia, Italy, 26-28 January 2011.

[3] M. Sichitiu, M. Kihl, “Inter-vehicle communication systems: a survey”, *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 88–105, 2nd quarter 2008.

[4] R. Uzcategui, G. Acosta-Marum, “Wave: a tutorial”, *IEEE Communications Magazine*, vol. 47, no. 5, pp. 106-133, May 2009.

[5] P. Muhlethaler, A. Laouiti, Y. Toor, “Comparison of flooding techniques for safety applications in VANETs”, *Proceedings of International Conference on Intelligent Transport Systems Telecommunications*, pp. 1–6, Sophia Antipolis, France, 6-8 June 2007.

[6] M. Torrent-Moreno, J. Mittag, P. Santi, H. Hartenstein, “Vehicle-to-vehicle communication: fair transmit power control for safety-critical information”, *IEEE Transactions on Vehicular Tech*nology, vol. 58, no. 7, pp. 3684-3703, September 2009.

[7] T. Nadeem, P. Shankar, L. Iftode, “A comparative study of data dissemination models for VANETs”, *Proceedings of International Conference on Mobile and Ubiquitous Systems*, pp. 1–10, San Jose, USA, 17-21 July 2006.

[8] S. Biswas, R. Tatchikou, F. Dion, “Vehicle-to-vehicle wireless communication protocols for enhancing highway traffic safety”, *IEEE Communications Magazine*, vol. 44, no. 1, pp. 74–82, January 2006.

[9] M. Nekovee, “Epidemic algorithms for reliable and efficient information dissemination in vehicular ad hoc networks”, *IET Intelligent Transport Systems*, vol. 3, no. 2, pp. 104–110, June 2009.

[10] K. Bür, M. Kihl, “Evaluation of selective broadcast algorithms for safety applications in vehicular ad hoc networks”, *International Journal of Vehicular Technology*, vol. 2011, Article ID 730895, 13 pages, February 2011.

[11] H. Lu, C. Poellabauer, “Balancing broadcast reliability and transmission range in VANETs”, *Proceedings of IEEE Vehicular Networking Conference*, pp 247-254, New Jersey, USA, December 2010.

[12] R. Chen, W. Jin, A. Regan, "Multi-hop broadcasting in vehicular ad hoc networks with shockwave traffic", *Proceedings of IEEE Consumer Communications and Networking Conference*, pp. 1-5, Las Vegas, USA, 9-12 January 2010.

[13] M. Artimy, W. Robertson, W. Phillips, “Assignment of dynamic transmission range based on estimation of vehicle density”, *Proceedings of ACM International Workshop on Vehicular Ad Hoc Networks*, pp. 40–48, Cologne, Germany, 2 September 2005.

[14] M. Sepulcre, J. Gozalvez, "Optimizing adaptive transmission policies for wireless vehicular communications", *Proceedings of IEEE Vehicular Technology Conference*, pp. 1-5, Calgary, Canada, 21-24 September 2008.

[15] The ns-3 network simulator, available online at http://www.nsnam.org , last accessed 8 August 2012.

[16] The network simulator ns-2, available online at http://www.isi.edu/nsnam/ns , last accessed 31 January 2012.

Biographies

Kaan Bür *earned his B.Sc. in control and computer engineering from Istanbul Technical University in 1995, his M.Sc. and Ph.D. in computer engineering from Boğaziçi University, Istanbul, in 1998 and 2006, respectively. From 2006 to 2007, he was a senior researcher at Faculty of Electrical Engineering and Computer Sciences, Berlin Institute of Technology. From 2008 to 2009, he was a research associate at Department of Electronic and Electrical Engineering, University College London. Since 2009, he is a researcher at Department of Electrical and Information Technology, Lund University. His current research interests include next generation broadband networks, wireless sensor and mesh networks, mobile and vehicular ad hoc networks.*

Maria Kihl *earned her M.Sc. in Computer Science and Engineering in 1993, and her Ph.D. in Communication Systems in 1999, both at Lund University. Since 2004, she has been an Associate Professor at the same university. During 2005-2006, she was a visiting researcher at North Carolina State University. Her main research area is performance of distributed telecommunication applications. Currently, she works on Internet and telecommunication server systems, vehicular networks, and multimedia applications in IP access networks..*