Virtual Access Points for Mobile Communication

This chapter presents a simple yet efficient dissemination algorithm called Virtual Access Points (VAPs). This work focuses on the problem of data dissemination in Infrastructure-to-Vehicle (I2V) and Vehicle-to-Vehicle (V2V) communication modes. The main objective of the technique, that can be used to spread public safety warning messages, is to extend the I2V network to areas where regular access points are not deployed. The technique is based on the DTN concept and the nodes exchange messages in an opportunistic way. When a vehicle moves near an Access Point, and receives a message, this vehicle becomes responsible for re-broadcasting it over the uncovered areas. This behavior is exemplified in Figure 4.1, where node A receives a message from the AP and afterward rebroadcasts it in a non-covered area.

4.1 Introduction

In the future pervasive wireless world, all roads and cities will be covered by roadside base stations and access will be provided to both pedestrians and vehicular users. However, for the moment, roadside units (RSUs), or Access Points (APs), are not always present, or may have been damaged as a result of a disaster. Furthermore, "historically, major disasters are the most intense generators of telecommunications traffic" [7]. The public communication

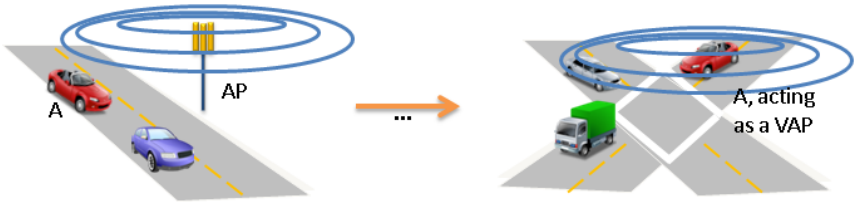

Figure 4.1: The VAP data dissemination technique

networks, even when available, may fail not only because of physical damage, but also as result of traffic overload. Therefore, the regular public networks alone are often not sufficient to allow rescue and relief operations [7].

For these reasons new and specific purpose mechanisms are required to ensure communication during catastrophic situations. With the occurrence of uncovered areas, the only possible communication mode is from one vehicle to another. This work relies on the existence of infrastructure-to-vehicle (I2V) and V2V communication to spread public safety messages among users over a defined region.

As introduced in Section 3.1, the next generation of cars will have radio capabilities. The method proposed here intends to take advantage of such capabilities to extend the coverage of emergency alert systems. Emergency warning messages are not frequent, but when they are issued they must be spread as fast as possible to all the people in the affected region. In this situation all the available means should be used to increase the awareness of the population regarding the imminent threat. We propose here that the available RSUs/APs act in partnership with regular vehicles to help on the spreading of messages that could be, for example, EAS warning messages in case of an emergency.

The proposed technique is called Virtual Access Points (VAPs) and it is a simple, yet powerful, technique to extend coverage to nodes outside covered areas. We consider a system like the one proposed by the RATCOM project [88], depicted in Figure 2.2. In the next generation of EAS, sensors will capture data and, if a real anomaly is detected, warning messages will be distributed automatically over the endangered region.

The RATCOM alert system is composed of two main components: one ascendant and one descendant. The ascendant component is responsible for sensing the related data, filtering false positives and retransmitting the relevant collected information to the coordination center. The descendant

component is responsible for spreading the information of the imminent dangerous situation among the authorities and population in general. This work focuses on this last phase: we try to increase the awareness of the general population of the imminent danger using the wireless medium and V2V communication.

4.2 Virtual Access Points for mobile communication

4.2.1 Protocol explanation

The main focus of the Virtual Access Point technique is to decrease the areas not covered by roadside APs so as to minimize the problem of intermittent access to mobile nodes. If we are able to decrease this problem, then even stream traffic for mobile users may be enabled. This work is based on opportunistic node contact. The proposed protocol prime for the simplicity as the duration of the contact opportunities between mobile nodes tends to be small. Chaintreau et al. points that for human mobility patterns the contact duration follows a heavy tailed distribution [27] [28]. They observed that the fast contacts are the most common ones among nodes in real world mobility patterns.

The protocol can be summarized as follows. Each node, after receiving a message, caches it and can thus later become a VAP, acting in a similar way to a relay node. Note however that, instead of just resending the messages, the VAP stores the message and may send it more than once or not at all depending on the caching strategy and depending on the locations has it passed by. VAPs strive to supplement the lack of real APs in a given area broadcasting messages received previously from other AP or even VAPs. A node acts as a VAP if it is neither in the range of an AP nor of a VAP and its distance from the nearest AP is $2r$, where r denotes the AP transmission range. This in practice means that a node is allowed to act as a VAP only when it is at a distance where its MAC layer does not detect any APs above a very low SNR and where it will not interfere with the signal of other APs. We also assume that the MAC layer takes care of solving conflicts and of treats the medium access problem. This application is just one of possibly many others running in the network: this is why the number of messages of the stream application is controlled.

In case a node senses another node acting as VAP in the same region, it gives up being a VAP, even if it lies in an area where it could act as

one. Therefore, the first node to broadcast VAP messages in a given region becomes the VAP for that time interval. Nodes are not allowed to act as VAPs during two consecutive time intervals. The high level VAP algorithm is presented in Algorithm 1.

Algorithm 1 - The VAP high level algorithm, from the point of view of a mobile node that can act as VAP

```
1: // At each time interval
2: if (Received a message && message is from an AP/VAP) then
3:     Stores the message in the cache;
4:     if (cache is full) then
5:         Throws away the oldest message;
6:         Stores the new message;
7:     end if
8: end if
9: // Verifies if will act as a VAP or not
10: if (node has messages in the cache && node is in a position where it could became
    a VAP && node was not a VAP in the last round && it did not receive any message
    from other VAPs this round) then
11:     Randomly chooses either to becomes a VAP or not;
12:     if (node become a VAP) then
13:         Chooses, from the cache, the message(s) to rebroadcast;
14:         Rebroadcasts the chosen message(s);
15:     end if
16: end if
```

Figure 4.2 shows a typical scenario where a vehicle A acts as VAP providing access to vehicle C. Vehicle D that is receiving a new message from the AP will also, at some point in the future, rebroadcast this received message to the nodes spread over the uncovered area. For all practical purposes we consider that there is no difference between the messages received from a road side AP or a VAP. The propagation mechanism is cooperative and transparent, from the point of view of the receiver. The system is a best effort one; there are no guarantees that every node will receive all stream packets, but using VAPs, we aim to increase the chances for timely reception.

Even in case of a severe catastrophe, or a huge terrorist attack it is unlikely that all the RSU's would break down at the same time. We consider that some RSU will be able to rebroadcast the warning message to the population. After that, the vehicles that received the warning will also be able to spread this information to the other vehicles on their path, which in their turn may do the same. As explained in Sections 3.3.1 and 3.3.1, this kind of propagation scheme is normally referred in the literature as epidemic and nodes act in a *store-carry-and-replicate* paradigm [99].

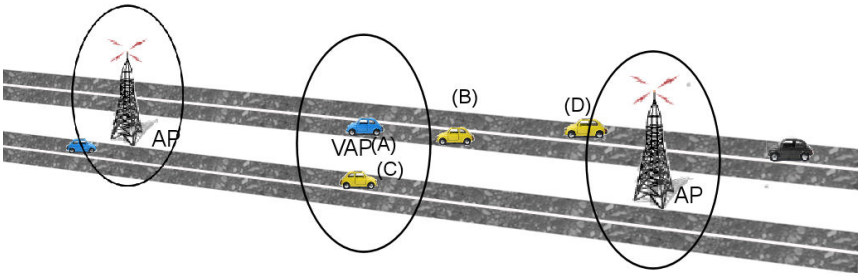

Figure 4.2: The VAP technique on a road coverage vision

To decrease the waste of resources and avoid medium access problems, vehicles act as VAPs only when they are out of the range of a real AP and if they have not received any communication from another VAP during this time slot. In the case of a disaster scenario, this kind of cooperative behavior may be the only way to disseminate useful and general information through the network.

4.2.2 Analysis

As we will see in Section 4.3 the technique successfully decreases the uncovered areas, but it has a cost. The cost can be measured in terms of the increase in the number of messages sent through the network. Consider the target message as a limited size stream being generated at a constant bit rate (CBR): this means that during each second n packets, from the total message size (η), are generated from a source and spread through all real APs. Each AP then is in charge of re-broadcasting the received message to the nodes in its area. Assuming that part of the message is transmitted from each antenna just once, the increase in the number of messages sent (im) is upper bounded by:

$$im = \alpha - (nVAP * \eta), \tag{4.1}$$

where α is the total number of exchanged messages and may be expressed as:

$$\alpha \leq \beta = (nVAP * \eta) * t, \tag{4.2}$$

where β is the maximum number of exchanged messages in each interval of time, $nVAP$ is the number of virtual roadside units, η is the size of the warning message and t is the time the warning message is propagated. The minimum possible number of packets in the network is given by the number of mobile stations in the region times the size of the message. I.e. each

vehicle received the complete warning message just one time. This would be possible, for example, if the whole area was covered by RSUs. However, with a distributed communication algorithm this value is hardly achievable. However, it is clear that the number and locations of the RSUs will greatly affect the system performance. The points where vehicles will act as VAPs are directly related to the deployment of the RSUs. Well deployed RSUs can provide faster and more efficient message spreading over the target region.

Using the technique described in [19] we formally verified the VAP protocol behavior prior to its simulation. Our aim was to verify whether the protocol is loop-free or not. Surprisingly we found a number of situations where loops may occur. For example, considering Figure 4.2, the simplest loop scenario occurs in the following case: node A, acting as VAP, transmits the message M_1 that is received by node B. Suppose node B is faster than node A and starts to act as a VAP at a point ahead in the road, it can transmit message M_1, which if received by node A would characterize a loop. For this reason messages need to be equipped with unique identifiers (IDs). Once the node A receives a duplicated message, identified by the ID, it discards it, thereby preventing the loop formation.

Another type of message loop may occur, and is in fact desirable even. Again, let us consider Figure 4.2; supposing that node A acts as a VAP in *lane* 1, the message M_1 sent can reach the node C, going in the opposite direction in the *lane* 2. At some point in the future node C starts to act as a VAP and retransmits the message M_1 that is received by the node D in *lane*1. If node D does not have the message, it is stored and will be retransmitted in the future in case node D becomes a VAP. However notice that this case is not a loop in the conventional sense, since the nodes involved are different. Another point to observe is that this kind of loop is even desirable since it helps in spreading messages over the region. The buffer favors newer messages, so older messages will be ignored and removed from it.

Even though the VAPs do not transmit when they find out that there is another VAP in the same area, depending on the MAC layer protocol used, concurrent transmissions and hidden/exposed nodes problems may also occur. Here we consider the existence of a MAC layer mechanism to handle this, e.g. scheduler for IEEE 802.16 networks or CSMA/CA for IEEE 802.11 networks. However, even if collisions occur, the worst impact will be a waste of bandwidth in a region that was not previously in use anyway.

We also found out that there may be nodes in the network that never take advantage of the VAPs technique. There is no guarantee the mobile nodes will receive all the messages needed to fill their buffers, or a node traverses

the entire path from one AP to the other without receiving any message from other VAPs. This will happen if the node is unfortunate enough not be inside the VAP range of other nodes acting as VAP, or when the node itself is acting as VAP for others, and thus is not receiving messages from other VAPs. These situations are more likely to occur in sparse networks.

4.3 Experiments

We now present the evaluations made to determine the impact of the VAP technique over spreading the message through the network. We have three different sets of experiments. The first one evaluates the application of the technique in stream based traffic, the second set of experiments shows the impact of VAPs over different disaster scenarios and the third set evaluates the impact of VAPs over a warning message distribution occurring over a bigger suburb-like area.

4.3.1 Environment

The simulations were programmed on top of the Sinalgo simulator [96], developed by the Distributed Computing Group at ETH Zurich. All the experiments were conducted using Linux Fedora Core release 6 in an Intel Xeon 1.86GHz machine with 16GB of RAM. The graphs are presented with a five percentile and a confidence interval of 99%. Each point is the result of the mean of at least 34 runs with different network configurations. The sizes of the target areas and the period of simulation vary according to the experiment. The APs positions are chosen randomly and the APs are static. In the typical set up, messages are spread by the available APs at a rate of 1 message per second of simulation and the same message is distributed simultaneously by all the available APs.

The scenarios follow a realistic mobility pattern generated with the VanetMobiSim [53] tool. All simulations keep the same basic configuration and only one of the parameters is varied: these include the number and the position of APs, the size of the available cache, the size of the message, the type of disaster scenario and the time at which the disaster occurred during the simulation.

4.3.2 Evaluated Disaster Scenarios

One of the main objectives of this work is to create techniques that can work even during severe conditions. Considering this, some experiments

were conducted to determine the resilience of the VAP technique in disaster situations. Here we evaluate the impact of two kinds of disaster scenario, the first one is when the network is damaged by natural causes and the second kind is when the network is damaged by sabotage, possibly as a result of terrorist attacks. The tested scenarios evaluate the behavior of regular nodes, before and after the catastrophe. The nodes are the same and follow a realistic movement patterns. This does not mean that we claim that movement patterns will be the same before and after an earthquake, for example . However in the absence of real meaningful data, and considering nodes will still be able to move, we chose to use realistic mobility patterns as a way to test the use of the VAPs to improve the connectivity of the remaining nodes. The natural disasters evaluated here are earthquake and flooding, whereas the sabotage scenarios are power outage and network random failures. These disaster scenarios were abstracted in the simulation as follows:

- Earthquake: The network starts with all the APs and mobile nodes running perfectly. However, at some point, 80% of the existing APs are randomly damaged and excluded from the network. This abstraction permits us to evaluate the effect of the technique when a major part of the APs disappear randomly from the network without any warning.
- Flooding: The evaluated scenario is a flash flooding [35] one. This kind of flooding is common in mountain regions in spring, heavy rainfall during the tropical rainy season and in the case of dam failures. This situation is abstracted in the simulations by the random disabling of a slice of 20%, either horizontally or vertically, of the middle section of the network. All the APs in this segment of the network are disabled. This is meant to simulate a river crossing the city that flooded in a sudden way.
- Power outage: In this scenario, we divided the evaluated scenario in four quadrants. During the simulation one of the four quadrants is randomly chosen and all APs in that quadrant are disabled. Complete blackouts are rare in developed countries, but power outages in cities are relatively common if some problem occurs in a specific power station, power line or other part of the distribution system. Commonly the effect of these failures is that part of the power grid goes down leaving part of the served region without energy. Such problems could occur by accident, or as a result of sabotage.
- Random network failure: In this scenario random network APs fail and disappear from the network during the regular network operation.

The degradation of the network coverage, in this case, is gradual, in contrast to what occurs in the other scenarios. This kind of generalized and chronic failure scenario could be triggered by hacker actions or physical sabotage of the nodes to deny access to the network.

4.3.3 Stream oriented traffic

This section presents the impact of the VAP technique over stream oriented traffic. We examine two types of environments: a highway segment and a city section. The road segment considered is $5Km$ long having four lanes, two in each direction with cars going back and forth on it. For the city environment we chose a $2km^2$ area of Washington D.C. city center with cars distributed through it. The area used for the experiments is depicted in Figure 4.3. For each scenario we have 40 different configurations of 10 simulation minutes, with 200 vehicles and a transmission range of $100m$. For the city environment the nodes minimum speed is $18km/h$ and the maximum is the maximum allowed on that specific road. For the road environment the vehicles minimum and maximum speeds are $60km/h$ and $110km/h$ respectively.

Figure 4.3: Map showing the Washington D.C. area we use for the simulations

All experiments keep the same basic configuration and a single parameter is varied in each one. The varied parameters are: the stream transmission rate, the number of static APs and the method VAPs use to select messages to re-broadcast. The source of the stream generates an "infinite" Constant Bit Rate (CBR) traffic from 1 to 3 messages per second. We call it "infinite" because the stream is constant and, for this set of experiments, does not repeat. It simulates a web radio broadcast, what could be an information channel news stream. Each generated scenario has a number of APs placed randomly. The number of APs tested for the city environments where 2, 25, 50 and 100. For the road environment the number of APs evaluated where 2, 5, 10 and 15. Every mobile node has a limited size buffer where it stores the last received messages. During cache replacement the oldest stream message, with lower stream ID, is discarded first, regardless of whether it was the first to be received or not. The three ways the VAPs messages are chosen to be rebroadcasted are random, oldest message first and newest message first.

The use of the VAP technique effectively allows us to increase the coverage of the real APs through the help of mobile nodes. These nodes coordinate to increase network coverage area. The graphs of Figures 4.4 and 4.5 demonstrate typical histograms of messages received in a 2Km simulated square of Washington DC and on a road segment, respectively. We can see that the mobile nodes can cache messages originated from the APs, and act as VAPs to other nodes in non-covered areas. Thus, the nodes, collaboratively, help to forward packets to areas previously uncovered and unused. The VAP technique was first designed to be used in road like environments, but as shown in Figure 4.5 metropolitan environments can also benefit from it. If we compare the plotted map with the actual area map, presented in Figure 4.3, we can even guess the roads and main intersections from it.

VAPs were first devised for road environments, however, as Figure 4.6 and Figure 4.7 show, it is valuable in both scenarios. Figure 4.6 demonstrates the behavior of the VAPs for a road environment displaying the number of unique messages received. Unique messages are defined as messages received by a mobile node for the first time. Numbers of unique messages start to decrease around the 200s because at this point the caches of the nodes start to saturate with stream messages and diversity of messages among the nodes caches decreases. This does not occur as much in the city environment as it does in rural environment of the simulated scenarios. In the road scenario the cars perpetually move along the two opposite highway directions. Thus, the nodes exchange more messages but of decreased diversity. In the city environment, however, nodes follow dissimilar paths which results in diverse cache contents. The number of lost messages decreased between 10% to 15%

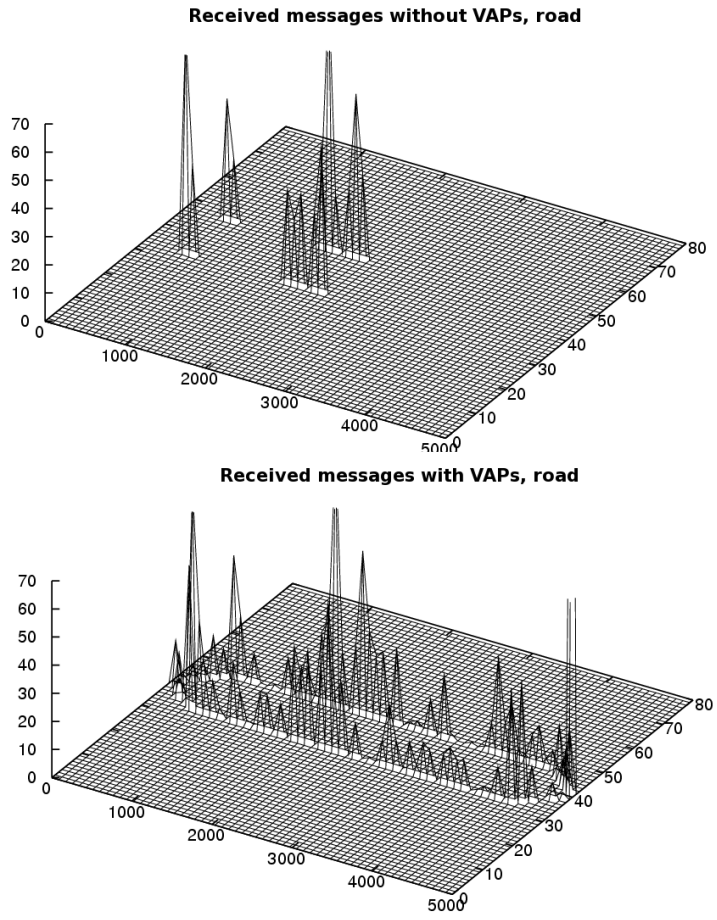

Figure 4.4: Typical histogram of the messages transmissions over the road segment observed

for city environments while for the road environment it decreased between 10% and 27.88%.

Figure 4.8 shows the difference of having 2 or 25 APs in the city scenario for varying bit rates. Both the number of APs and the bit rate influence the number of unique messages received in total. However, as expected, the number of unique messages for scenarios where VAPs are not present is nearly constant, as it only depends on the nodes passing near the APs. Even when the bit rate increases we do not observe a significant increase in the number of unique received messages. When bit rates are increased from 1 to 3 packets per second, in the 2 APs case, the result is marginal. When VAPs

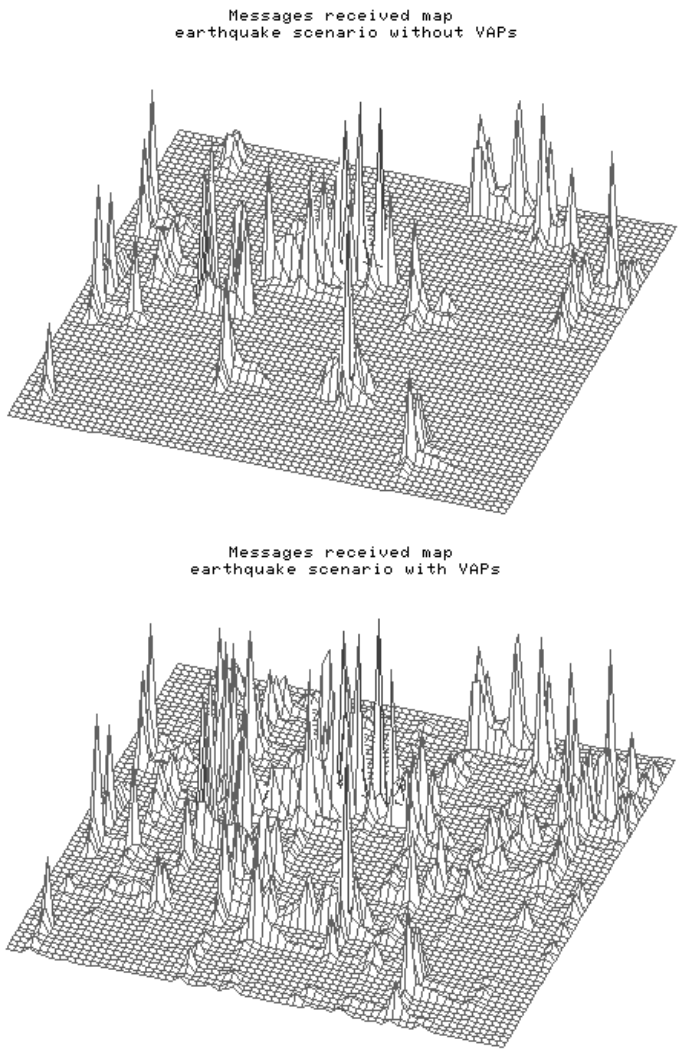

Figure 4.5: Typical histogram of the messages transmissions over the city experiment

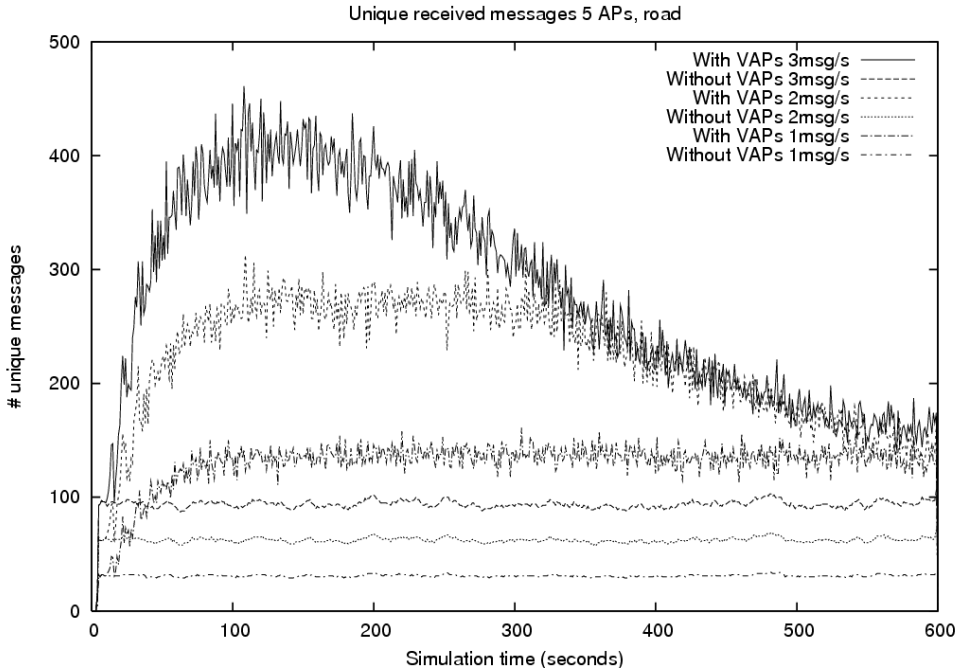

Figure 4.6: Unique received messages through the 10 mins of simulation for the road environment with different traffic rates

are enabled, the number of unique messages received significantly increases, because 2 antennas are not enough to spread the information through the entire network. The VAPs take advantage of nodes caches to propagate messages which were previously lost.

However, the larger the covered area the lower is the gain the VAP technique presents. This becomes apparent when we look at the graph of Figure 4.9. The graph shows, for the same experiments, the messages first received through APs and VAPs. As the number of APs nodes increases in the road environment, the number of messages first received through VAPs decreases. The behavior is similar for the city environment.

The use of VAPs accounts for an increase between 61.7% and 134.57% on the total traffic of the network. However, since this increase occurs only in non-covered areas, it is not creating interference or delaying the system's APs. Nevertheless, evaluating the number of repeated messages is interesting. On Figure 4.10, the number of repeated messages for the networks that use VAPs and the ones that do not use it follows the same shape. Increasing the number of messages generated by the VAPs results in more repeated

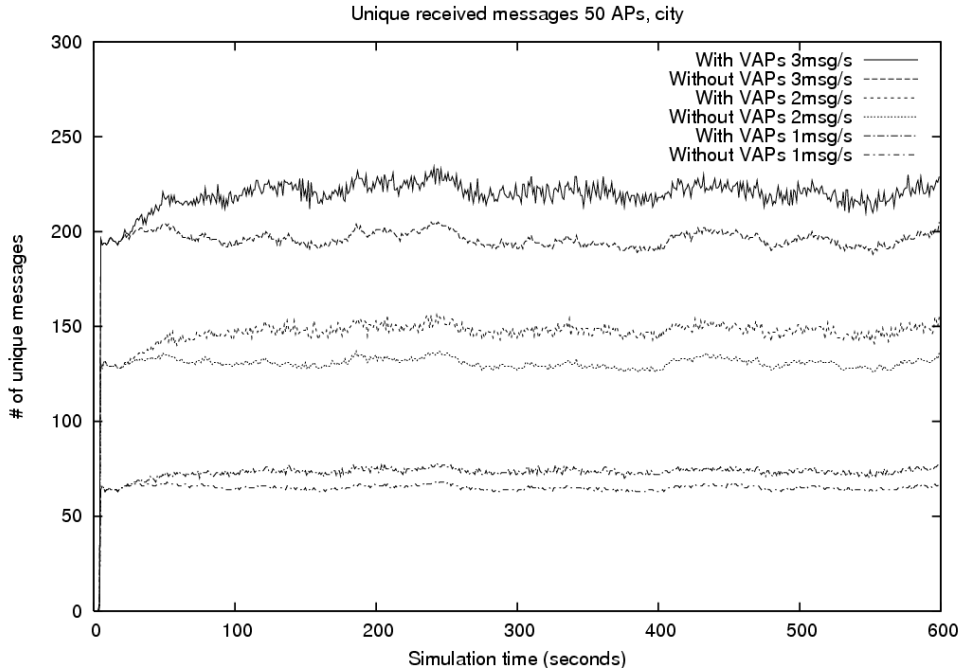

Figure 4.7: Unique received messages through the 10 mins of simulation for the city environments

messages. Figure 4.10 presents the results for different stream rates, as well as for different transmission rates for the VAPs. Each VAP node can either transmit at the same rate the stream generated or at 4 times this rate. For example, if the stream is generated at the rate of 1 message/second (m/s), the VAP can transmit cache messages either at $1m/s$ or at $4m/s$. The number of repeated messages increases based on the number of VAPs, but as the VAPs assignment is dynamic it decreases when the network coverage increases. This way the number of repeated messages also decreases, as there are less VAPs active. Ten is nearly the best number of APs for this scenario. Given less than 10 nodes, we have a lot of uncovered areas and with more than 10 the network gets so overprovisioned that APs start to interfere with each other and the number of repeated messages increase again, not because of the VAPs, but because one mobile node starts to receive messages from more than one AP.

Regarding the VAPs message spreading policies of random, older to newer and newer to older, all three presented nearly the same results. However, on

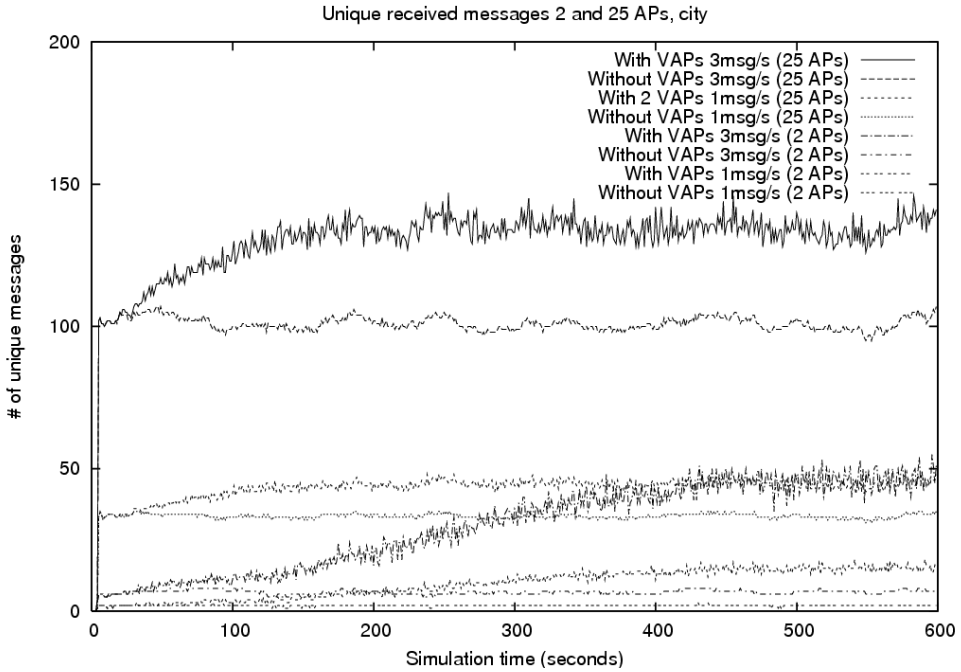

Figure 4.8: Unique received messages through the 10 minutes of simulation for the road environment with different number of APs and traffic rates

average, the random policy, i.e. the VAP node sending a random message from the cache, performed slightly better than the others.

4.3.4 Disaster resilience for stream traffic

For this set of experiments we used the same city area described in the previous set of experiments. A $2km^2$ area of the Washington D.C. city center with cars distributed over it. For each scenario we have 40 different configurations of 30 simulation minutes, with 200 vehicles and a transmission range of $120m$. For the city the nodes' minimum speed is $18km/h$ and the maximum is the maximum allowed on that specific road based on the data provided by the Topologically Integrated Geographic Encoding and Referencing (TIGER) system of U.S. Census Bureau. Each generated scenario has a number of APs placed randomly. The nodes are initially spread uniformly over the roads of the observed area and then follow the VanetMobiSim realistic mobility model. All experiments keep the same basic configuration but the number of and the locations of APs are random. On average we allocate

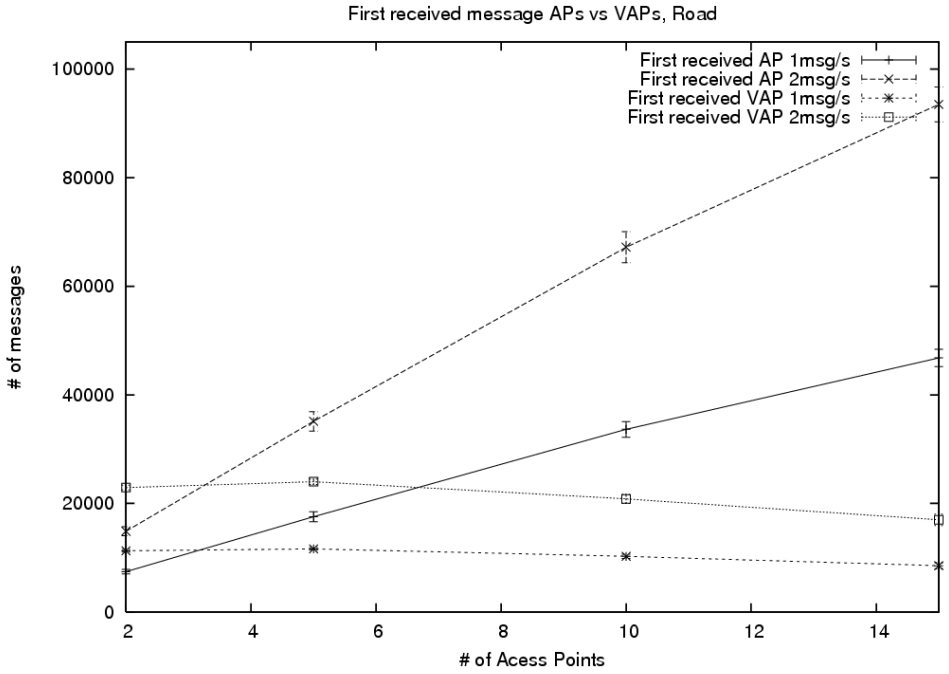

Figure 4.9: Number of messages first received from an AP and VAP

40 APs but the number varies up to 100. The source of the stream generates CBR traffic of 1 message per second, distributed simultaneously by all the available APs. We vary the number of APs, size of the cache, disaster scenario and time, during the simulation, when the disaster occurred.

Figure 4.11 shows the influence of the initial number of APs in the network on the percentage of messages received: for these simulations, the disaster occurs at the beginning of the simulation, and therefore the initial number of APs represents the number before the disaster takes place. We can observe that for all cases the VAP technique provides an increase in the number of stream messages received. The percentage of the stream traffic received is affected by the initial number of APs in the network, with larger number of APs, causing more extensive spread of information in the network. This makes the VAPs efficient for local traffic dissemination. In the best case, when no failure occurred in the network and all nodes work perfectly, using VAP technique provides an increase in the number of received messages that ranges from more than 700%, when the number of APs is two, to 16.6% when the initial number of access point is 100. Note that this ratio

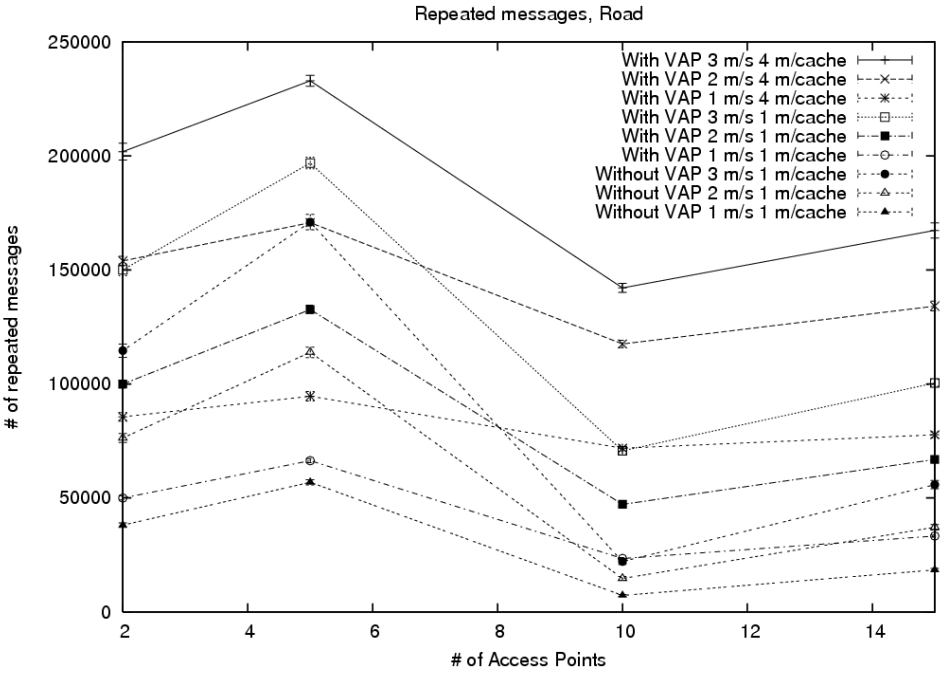

Figure 4.10: Repeated messages for the road environment

is caused by the fact that VAPs allow us to maintain communications in cases where otherwise the system would collapse.

One hundred access points represent, on average, a coverage of 58% of the total simulated area. As expected, the gain diminishes as the space covered by access points increases. This occurs because the VAPs are well behaved and, as it is an opportunistic protocol, the nodes act as VAP only when they are outside the range of any AP and any other VAP. With the increase of the network coverage by the real APs, the regions where a node could act as a VAP decrease and, therefore, the number of messages received through the VAPs decrease. For the disaster scenarios we can detect the same general behavior. Consequently, the percentage of the received streams is larger when the initial number of nodes increases. However, the corresponding gain introduced by the VAPs decreases. For the earthquake scenario the gain varies from 1615.91% to 71.95%. In this scenario 80% of the network is damaged in the beginning of the simulation, which explains the enormous gain. In this scenario, the number of actual APs is really small and almost all the delivered messages are done through VAPs. We call gain the percent

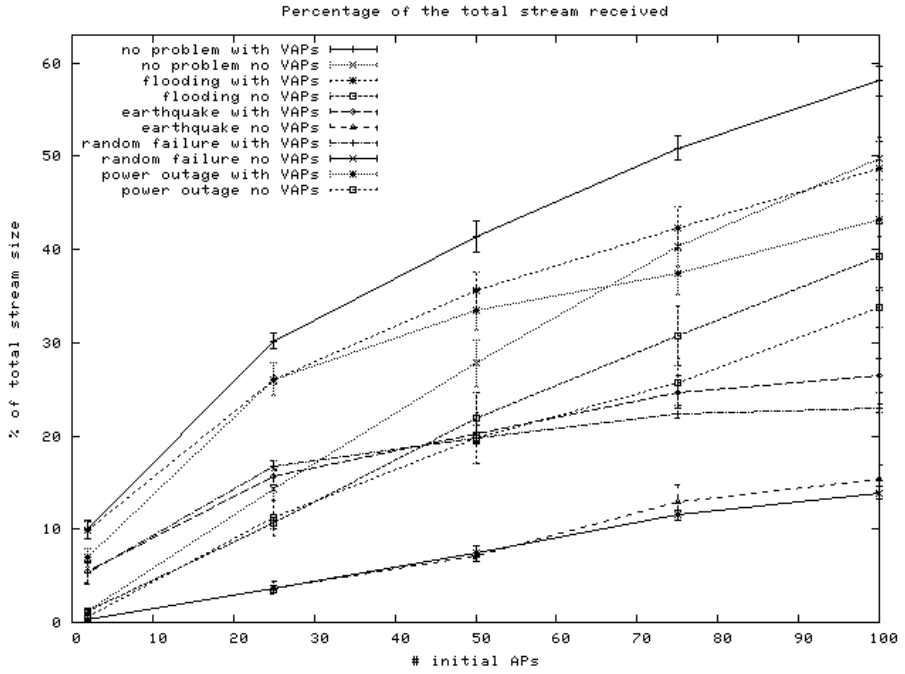

Figure 4.11: Average percentage of messages received in the network as a function of the initial number of APs for the evaluated disaster scenarios

of traffic delivered with help of VAPs over the amount initially delivered without the use of the technique. For example, if we double the number of delivered messages we say the gain attributed to the VAP technique is 100%. For the flooding scenario the gain varies from 753.88% to 24.27%. In the power outage scenario case, the gain varies between 1122.05% and 28.08%. As we can see, the gain reflects the fraction of the initial network affected by the disaster, in the flooding scenario 20% of the network is damaged and for the power outage one fourth of the network is affected. The larger the damage in the network, the more relevant the traffic received through VAPs.

For the random network failure scenario, the intervals between failures are random, distributed uniformly throughout the simulation time. By the end of the simulation only a few nodes remain functional. The damage for this scenario is not huge at first, unlike in the earthquake scenario. However the damage is constant through time. So that, by the end of the simulation, the damage caused to the network is comparable to that in the earthquake scenario. Figure 4.15 shows the behavior of the random failure and earthquake as a function of time. For the random failure, the gain varies from 1585.35% to 65.18%, close to the values estimated in the earthquake set-up.

We additionally vary the buffer size and the time the disaster occurred in the simulation. The results are basically equivalent; the only difference is a small increase in the total number of received messages, when we delay the disaster start time.

The graphs in Figures 4.12 and 4.13 show the number of duplicated messages received by the nodes during the experiments as a result of the application of the VAP technique versus the disaster start time and the size of the cache, respectively. The values for both graphs are relatively stable. This means that the duplicate messages have a low correlation with the size of the cache and the time the disasters started. As we can see in Figure 4.12 the biggest VAP overhead is around 32% of the total number of messages sent in the stream. However, on average 10% of the duplication results from nodes receiving duplicated messages from APs. So the real overhead caused, in the worst case, for the VAPs is about 22% of the network stream traffic generated. When there is no disaster, all 100 APs are working without any problem and a larger part of the stream is received by the mobile nodes from the APs. However, for the disaster scenarios, where not all the APs inject traffic into the network, the overhead varies between 9% and 16% for the flooding scenario, 16% and 19% for the random failure scenario, 24% to 26% for the flooding scenario and between 22% and 25% for the power outage scenario.

In Figure 4.14 we can observe that the number of messages received per cycle increases when we use VAPs. Using VAPs, the variability of that range increases. This is to be expected since VAPs is a best effort mechanism, and not as effective as APs would have been if they were available.

Figure 4.15 shows the number of unique messages received over time. Clearly, the use of VAPs increases the number of unique messages consistently over time. In this graph it is interesting to notice the behavior of the random failure scenario compared to the no disaster and earthquake ones. In the beginning of the simulation the random failure and the no failure results closely resemble one another. However, as time passes, the network degrades gradually in the case of the random failure scenario. VAPs decrease the impact of the APs failures, and enable mobile nodes to receive new messages even when virtually no mobile node receives messages directly from the APs. During the simulations it is guaranteed that, for any scenario, at least one AP exists and broadcasts new messages. If no VAP existed, only nodes in range of this AP would receive these new messages. Using the VAP technique these few nodes may spread the new message throughout the network.

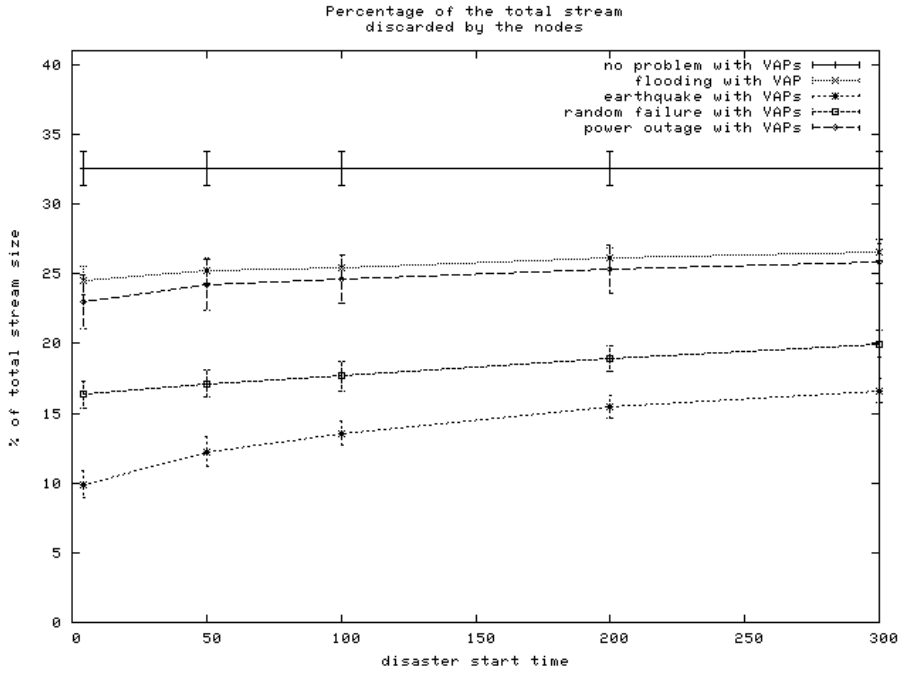

Figure 4.12: Number of duplicated messages received as a function of the initial time of each disaster

4.3.5 Warning message propagation

This set of experiments observes the use of VAPs in a suburb area. The evaluations were carried out in a $15 \times 9km^2$ area that encloses Sophia-Antipolis in the south of France, as depicted in Figure 4.16. The simulations were conducted with 1000 nodes with 200 meters communication range and speeds varying between $40km/h$ and $90km/h$. The nodes arrive randomly and are placed uniformly over the observed area.

We vary both the number of APs and the size of the messages, and analyze the impact of the occurrence of different disasters over the VAPs performance. The source of the stream generates CBR traffic of one packet per second that is distributed simultaneously by all the available VAPs. If the warning message is too big to send in one time interval, it is divided into smaller packets and these are broadcasted, one packet per second, continuously in a cyclical way. We consider transmission intervals of one second. Notice that this set-up is different from the stream-based one. Here we consider smaller messages and the same message is repeated.

The graph in Figure 4.17 shows the number of nodes that received the

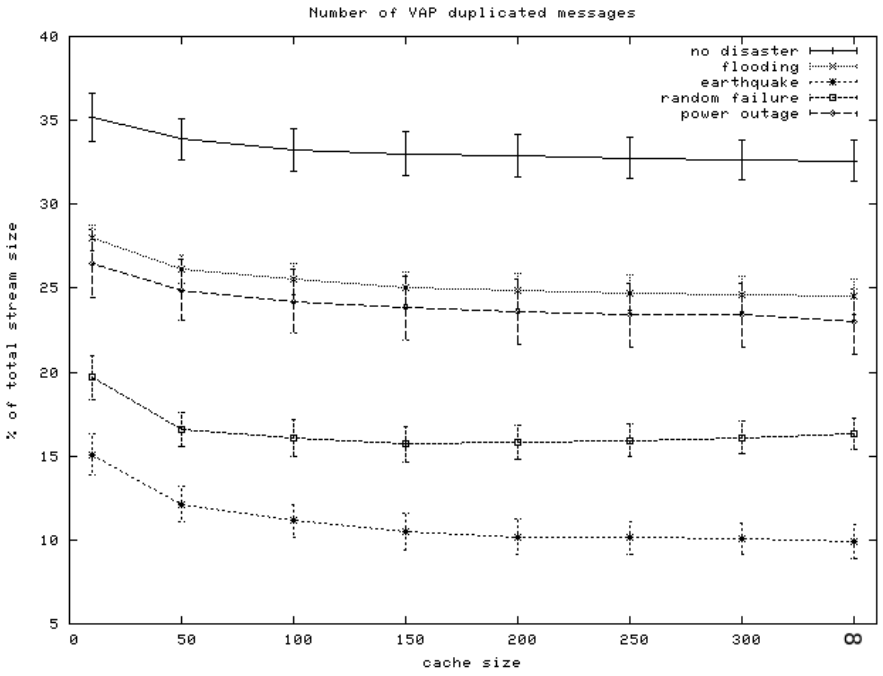

Figure 4.13: Number of duplicated messages as a function of the size of the cache with 100 nodes and the disasters occurring in the beginning of the simulation

one packet warning message for the different disaster scenarios. For all the scenarios evaluated with 10 initial VAPs, the use of VAP enabled the distribution of warning messages to all the network nodes. The most severe disaster evaluated is the earthquake one. In this scenario 80% of the initial VAPs are damaged during the experiment. However, even in this situation the VAPs delivered the warning to all the nodes in the region in less than 20 minutes. Even though the mechanism used to decrease the number of RSUs is different, for the earthquake and the random failure scenarios, their results are close. This occurs because with time the number of damaged stations in the random failure scenario increases. At the end of the simulation the number of remaining RSUs is nearly the same for both scenarios, however the smoother degradation of the random failure scenario grants it a better performance, when compared to the earthquake one. For the flooding scenario, only the nodes on the central strip of the area are removed. Although this affects the total number of nodes that received the warning message completely and slightly increases the time required to distribute the message to all the nodes in the network, the vehicle movement compensates for the

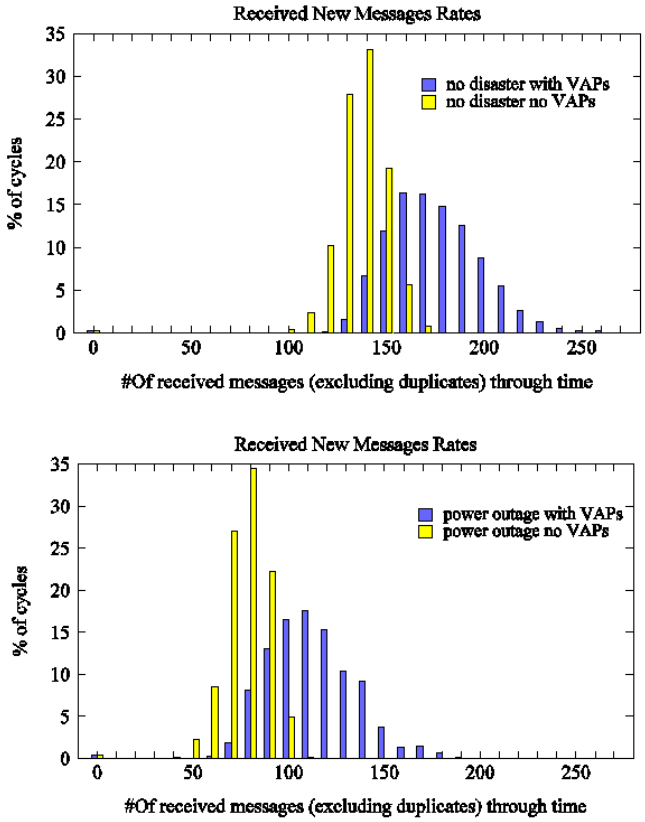

Figure 4.14: Transmission rate and variability

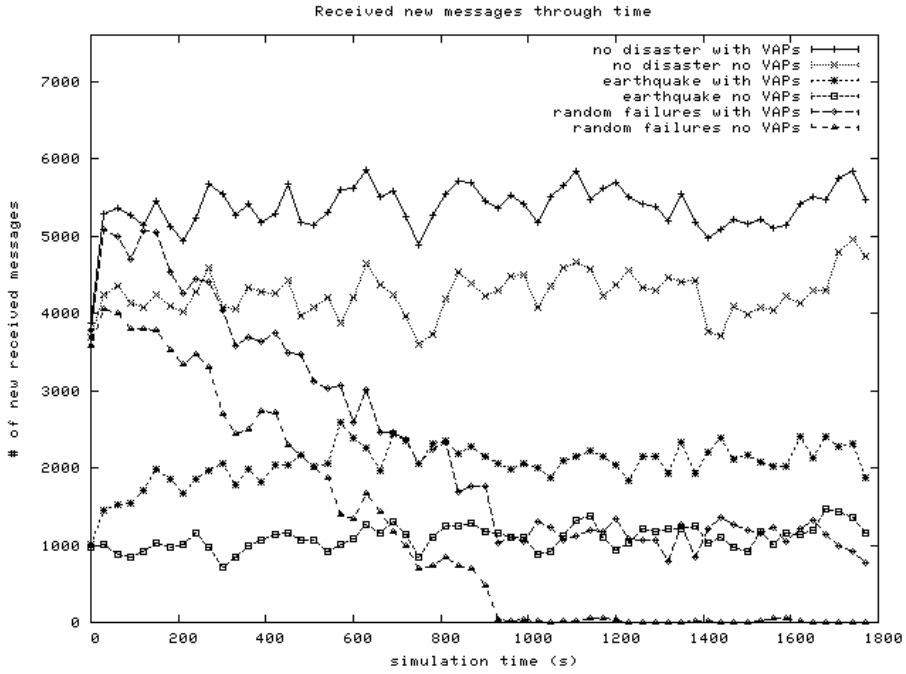

Figure 4.15: Received unique messages over time for the no disaster, the earthquake and the random failure scenarios

lack of RSUs in the central part of the area: nodes that did not receive the message because they were in that region, may later move to a region that is covered by RSUs.

We can perceive in Figure 4.17 that when no disaster occurred, the number of nodes warned is nearly 100%, regardless of whether VAPs are used or not. Indeed, the final number of nodes aware of the message is similar, when we do not consider any disaster. However the graph of Figure 4.18 shows the time it takes for all the target nodes to receive the message. We consider transmission cycles of one message per second, i.e. each second the warning message, or a part of it, is broadcasted. The plot shows the time when all nodes in the network received the warning message. Whether all nodes had received the messages or not the simulation experiment stops after 3600 seconds. If any node failed to receive the message within that interval, the registered time is 3600 seconds. Without the use of VAPs the network needs more than 200 APs to be able to spread the message to all the nodes in less than one hour. With the use of the VAPs, even in the worst case scenario of an earthquake with only two functional VAPs remaining, it takes around

Figure 4.16: The area in Sophia Antipolis used for this set of tests

20 minutes to spread the warning message over all the nodes in the region.

The tendency is that the time required to spread the warning message decreases when the number of VAPs increases. However, the gains become comparatively smaller when number of VAPs increases beyond 50. If we consider the no disaster scenario, if we increase the number of RSUs from 10 to 50 we speed up the message distribution by 28.8%. However, when we increase the number of RSUs from 50 to 500 the gain is 29.8%. I.e. with 50 VAPs we are able to warn the whole population in 8 minutes, whereas if we increase the number of RSUs to 500, the process will take around 5 minutes. This result is interesting since it shows that the increase in the number of RSUs does not linearly impact the time needed to warn the population over a given target area. This means that we could decrease the number of RSUs, and the cost of the system deployment, without compromising significantly the quality of the service offered. This effect is also clear from the graph of Figure 4.19. From this graph we see that when we increase the number of RSUs we do not increase proportionally the number of nodes that receive the warning message. Even without the use of VAPs, the node coverage

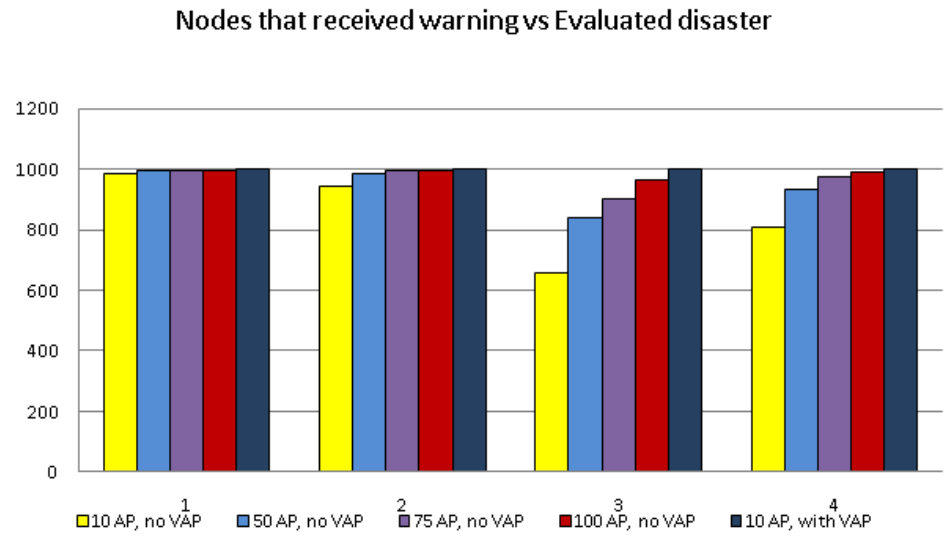

Figure 4.17: Number of nodes that received the warning message taking into account the evaluated disaster scenario

for all scenarios, except for the earthquake one, is almost 100% with only 50 RSUs. However, this value of active RSUs also holds for the earthquake scenario. In the earthquake scenario almost all the nodes are warned when we increase the number of initial VAPs to 200, i.e. when on average 40 RSUs are working during all the experiment: this is roughly the same number of nodes as in the other scenarios.

The apparent discrepancy between the graphs of Figure 4.18 and Figure 4.19 is given by only a small percentage of vehicles that did not receive the warning message during the simulation time. Because of their mobility patterns these nodes did not cross any VAPs during all the evaluated time. When we use VAPs we increase the coverage of the EAS, which permits not only to reach these nodes, but to reach them faster.

The graph of Figure 4.20 shows the percentages of messages first received through VAPs and of those received through real APs. The percentages on the graph are for the one packet size warning message and no disaster scenario. As expected when the number of APs increases the percentage of

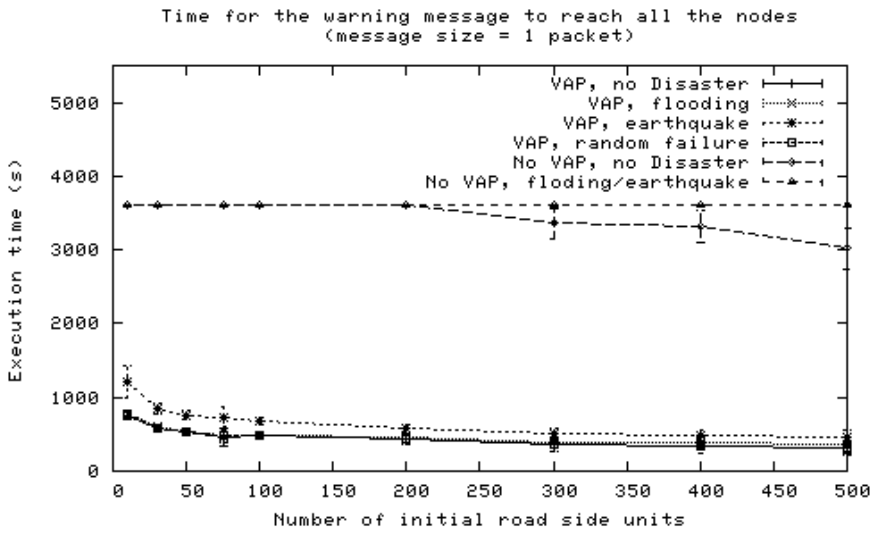

Figure 4.18: Average time for the warning message to reach all the nodes in the region. The simulation stops after 3600 seconds, this means that scenarios that had their time registered at 3600 seconds did not deliver the message to all nodes

packets delivered through VAPs decreases. Vehicles when acting as VAPs are really well behaved, if they perceive the presence of an AP or another VAP they defer retransmitting the warning messages. When we have 10 RSUs the percentage of roads covered by the VAPs is around 3%; on the other hand, when we have 500 RSUs spread randomly throughout the target area the percentage of roads covered by these RSUs is nearly 70%. This is roughly the same percentage of nodes that received the message through VAPs in the graph of Figure 4.20. It is clear that in the extreme case, if we had 100% of coverage, the VAPs would not increase the number of distributed messages. However, not only is it extremely expensive to have 100% of coverage, but also in the case of a disaster, the deployed infrastructure could be severely damaged. The main advantages of VAPs are their dynamicity and capacity to reach non covered areas.

The graph of Figure 4.21 shows the number of nodes that have received the whole message for increasing warning message sizes. As anticipated, increasing the size of the message decreases the number of nodes that receive it completely. However, the use of VAP provides an increase in the number

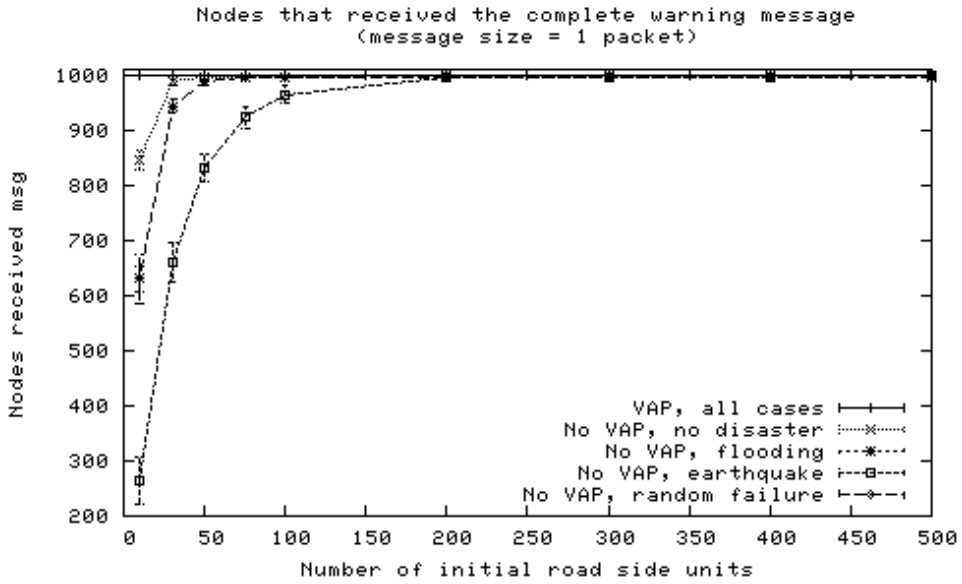

Figure 4.19: Number of nodes that received the warning message versus the number of roadside units on the network

of nodes that do so; comparing to the case without VAPs this increase varies from 14.4% to 60.8%, thus leading to a relatively stable number of warned nodes, even with the increase in the size of the message.

The experiments show that the proposed method increases the coverage and decreases the time required for all the nodes in the network to receive the message, however this has a cost. One of the ways to measure this cost is counting the number of repeated messages received by the nodes. The graph of Figure 4.22 shows the average number of repeated messages received by the nodes. The number of duplicated messages is considerably bigger when we use VAPs. The augmentation in the number of messages is also expected since the algorithm is an epidemic one. However, it is important to call attention to the fact that this traffic occurs in areas that had no communication before.

The number of duplicated messages, observed in the Figure 4.22, decreases when we increase the number of RSUs. This behavior is linked to the results observed in the graph of Figure 4.20. Again, when the area covered by the RSUs increases the areas where vehicles may act as VAPs

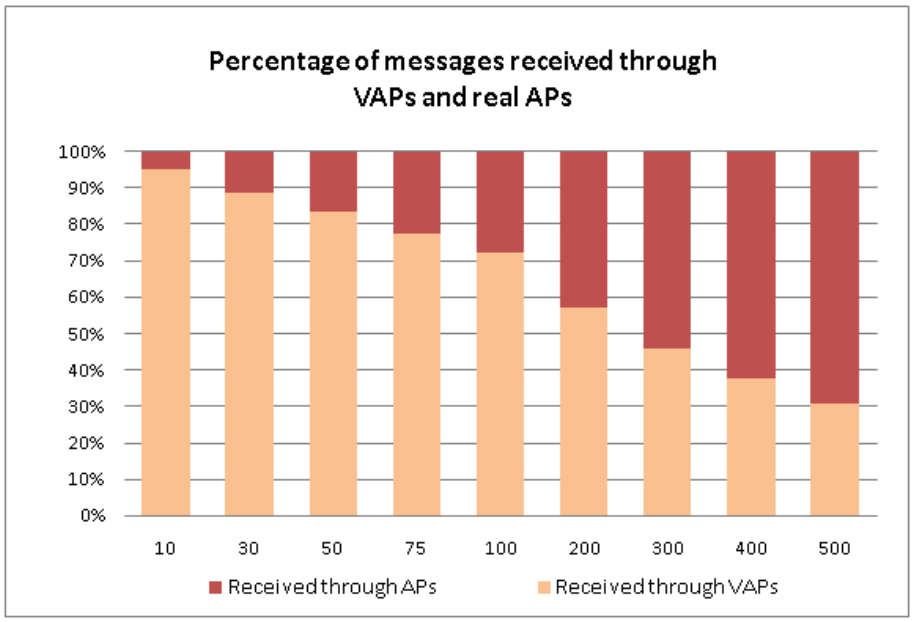

Figure 4.20: Comparison of the percentage of received messages through virtual roadside units and real road side units for one packet warning message vs. the number of road side units

decreases. From the graph in Figure 4.22 we can also observe that, apart from the earthquake scenario, the amount of traffic generated over the different scenarios does not vary significantly. As we can see in formula 4.2 the overhead is a function of the number of VAPs not RSUs. The earthquake scenario is a particular case, especially for small numbers of initial RSUs, for two reasons. First because after the disaster the number of APs is extremely small, so the area where vehicles may act as VAPs is bigger. The second factor is the small diversity of routes, when we have smaller number of APs. A vehicle only starts generating traffic after receiving the first message. When we have a small number of APs the number of sources of traffic is low, and the amount of routes nearby these APs is smaller. Nodes have then more chance of sending the message to nodes that have already received it. The nodes that really need to receive the message are the ones more distant from the AP.

The behavior of the message propagation is similar to the wave generated when we throw a stone in a lake. The wave propagates in every direction, but it takes some time to spread through all the lake and reach its borders. The warning message spreads in a similar way, reaching new nodes at each step.

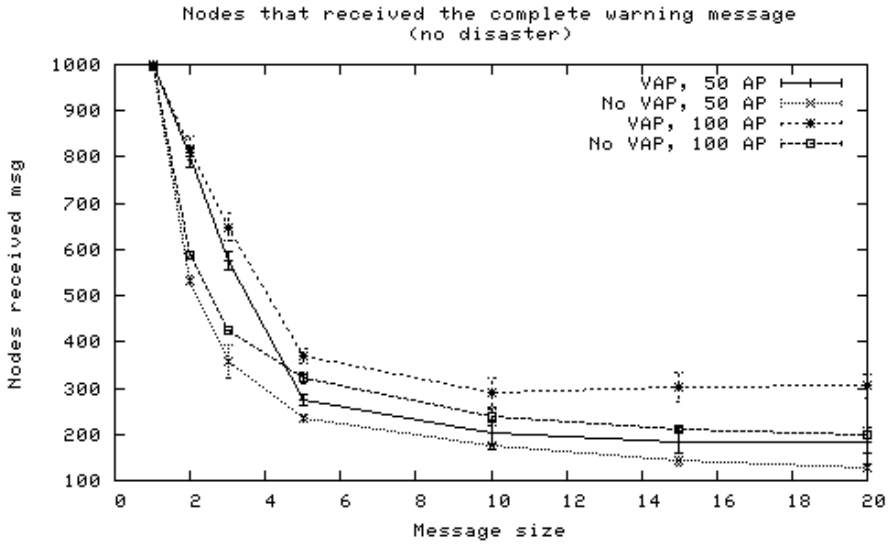

Figure 4.21: Number of nodes that received the complete warning message, vs. the size of the message

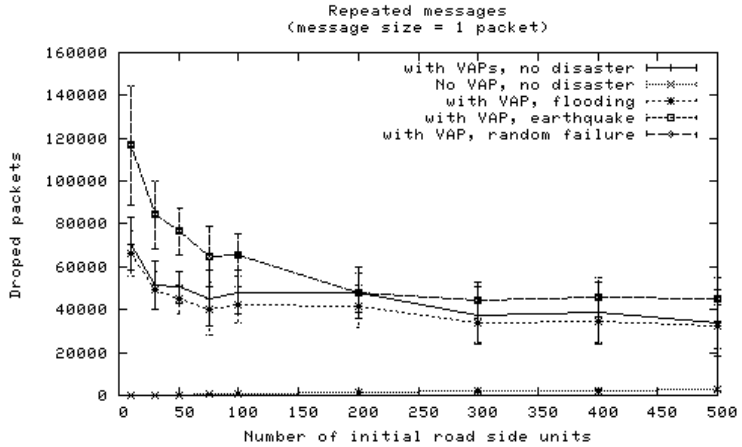

Figure 4.22: Number of repeated messages received by the mobile nodes during the simulation

If the number of VAPs is small the message wave takes more time to reach all the nodes in the network, as we can notice in Figure 4.18. The increase in the simulation time also leads to an increase in the number of messages received. However, when the number of VAPs increases the earthquake scenario starts to present a behavior similar to the one of the other disaster scenarios. None of the other scenarios presents such a severe loss in terms of APs. Even the random failure, which in the end loses a similar amount of APs as the earthquake scenario, does it in a gradual way. For the random failure in the beginning the number of RSUs is bigger and this increases the number of cars with the information. Consequently this increases the variety of the paths followed by the vehicles. This does not occurs in the earth quake where all nodes disappear at the same time.

4.4 Conclusions

The Virtual Access Point technique was proposed to increase network coverage for stream based and warning message traffic in normal situations and disaster scenarios. VAP is a simple yet effective method to increase network coverage for non-real time traffic. As discussed in the simulation results, we observe that the number of received messages for all the evaluated scenarios is increased, often impressively so, which is justified since our system manages to remain operational after the initial system has collapsed. Emergency alert messages are not frequent, but when issued they should be distributed as fast as possible to everyone in the affected region. Lives may depend on how fast and how broad the warning message was distributed.

The experiments show that the gain in the number of received messages may vary from nearly 1615.91% to 24.27% depending on the disaster scenario evaluated. Moreover, VAP is a valuable technique to disseminate network traffic even when no disaster has occurred and can operate transparently to the system. The gain resulting from the application of the VAP technique in the regular network scenario varies from 755.52% to 16.6%, depending on the number of APs disseminating data in the network. Since VAPs are only used in uncovered areas, the gain observed is negatively correlated to the AP coverage of an area. As a result of applying the VAP technique, the number of duplicated and irrelevant messages received by the nodes is increased. However, this traffic occurs in uncovered areas where it causes very low if any interference, and in the worst case, average traffic overhead is increased by approximately 27% in the experiments discussed.

We have also shown that even with a small amount of real access points,

using the VAP concept can broaden and significantly speed up the warning message distribution process. Our experiments show that even in severe conditions warning messages can reach all the observed nodes within a reasonable amount of time. On average, sending one packet per second the message was able to reach all nodes in the observed region, 15x9km^2, in six to seven minutes.

Part III

Crisis Management Phase Support

Chapter 5

Mesh Networks

Communication backbones in PSNs are normally organized as a mesh network. However, Wireless Mesh Networks (WMN) generally require self-organization and topology control algorithms to enable its broad use [6]. Studying this problem, of self organization in Mesh networks, has lead us to the main contribution of this part of the thesis, a topology control algorithm to maintain, in an efficient way, the topology of PSNs. Note that this technique can also be applied in ad hoc networks.

Because the next part of the thesis rely on WMN concepts, before moving on to the proposed topology management algorithm, this Chapter provides an introduction to WMN. We use the IEEE 802.16 standard [2] as an example, which could also be used to organize nodes in PSN. This chapter focuses on the main characteristics of WMN, how the nodes can communicate, using the IEEE 802.16 scheduling, to avoid inter-node interference and the possible types of services that can be provided over a WMN. These aspects are relevant to PSNs since these networks must be stable, predictable and provide QoS.

5.1 Introduction

Wireless Mesh Networks have been attracting a huge amount of attention from both academia and industry in the last few years. Indeed, WMN is now emerging as a promising technology for broadband wireless access [6]