

Intrusion Tolerance in Wireless Environments: An Experimental Evaluation

Henrique Moniz, Nuno Ferreira Neves, Miguel Correia, António Casimiro, Paulo Veríssimo
University of Lisboa*
{hmoniz, nuno, mpc, casim, pjv}@di.fc.ul.pt

Abstract

This paper presents a study on the performance of intrusion-tolerant protocols in wireless LANs. The protocols are evaluated in several different environmental settings, and also within the context of a car platooning application for distributed cruise control. The experimental evaluation reveals how performance is affected by the various environmental parameters such as the wireless standard, group size, and network topology. The distributed cruise control application demonstrates the practicability of such protocols, even when subjected to malicious faults.

1. Introduction

Intrusion tolerant protocols, which guarantee correct system behavior even if some of its components fail in arbitrary ways [11, 19], have been used in the past for several applications, including reliable communication, consensus and voting, and state machine replication [17, 14, 5, 4]. A few of these protocols have been implemented, however, their evaluation has been limited to wired LANs and WANs. The system models assumed by the protocols indicate that they were not designed with wireless communication in mind (e.g., they assume that nodes are fully-connected). Therefore, very little is known about the behavior of the protocols in wireless environments. Namely, it is unknown how well they will adapt to the distinctive characteristics of wireless networks, such as potentially lower bandwidths, higher failure rates and increased contention in the communication medium. Additionally, the kinds of devices that are many times utilized in these environments also have distinctive characteristics, for instance limited power supply and lower CPU capabilities, which most probably will impact on the performance of the protocols.

In this paper, we aim at understanding how well existing protocols can support the execution of intrusion-tolerant ap-

plications in wireless LANs. To contextualize our research, we have considered a particular application with demanding requirements – a *car platooning* service for *distributed cruise control*. In simple terms, the service is responsible for exchanging information in order to automatically control a group of vehicles traveling close together to some common destination (e.g., a set of trucks transporting a load across country). This application is quite interesting because it has some requisites both in terms of time and security: for example, velocity should be adjusted within a reasonable interval if a car wants to increase or decrease speed; and, correct behavior needs to be preserved even if external or internal attacks occur.

We developed an algorithm which implements the *speed agreement* operations of the distributed cruise control application. This algorithm is built using the services of an existing stack of intrusion-tolerant protocols for group communication – RITAS [16]. The services provided by RITAS are particularly interesting because they tolerate arbitrary (or Byzantine) failures, and they are built for an asynchronous system model. This synchrony model avoids all time dependencies, making the protocols immune to attacks in the domain of time (e.g., attacks which artificially delay some parts of the system). Moreover, it is well-suited for the considered environment because under certain conditions wireless networks are essentially untimely (e.g., noisy environments). RITAS offers several broadcast and consensus protocols, and copes with the FLP impossibility result for asynchronous systems by resorting to randomization [10].

Both the distributed control application and its supporting protocols are evaluated under several environmental settings. With the intention of not losing sight of the broader goal of the paper (i.e., to measure the feasibility of intrusion tolerance in wireless networks), the experiments were made using hardware with capabilities similar to what is expected to be found in automobiles (low-end personal computers), and with hardware that is typically utilized in wireless networks in general (PDA devices). For this purpose, a subset of the RITAS protocols had to be ported from Linux to the Windows Mobile platform. This task ended up being more

*This work was partially supported by the EU through project IST-FP6-STREP-26979 (HIDENETS) and NoE IST-4-026764-NOE (RESIST), and by the FCT through project POSI/EIA/60334/2004 (RITAS) and the Large-Scale Informatic Systems Laboratory (LASIGE)

difficult than expected due to differences in the thread management of Windows Mobile, and because of the limited support for some of the cryptographic operations.

To the best of our knowledge, there is no other study on intrusion-tolerant agreement protocols in wireless networks. There has been an important track of research on routing protocols for wireless ad-hoc networks that can be considered to be intrusion-tolerant, since they aim to tolerate routing misbehavior of some of the nodes. Two surveys are in [13, 1]. There is also one work on intrusion-tolerant broadcast in that kind of networks [9]. More recently, there has been research on crash fault-tolerant consensus protocols for wireless networks, presenting conditions for solvability, tradeoffs and efficient algorithms [6, 7, 20, 2, 12].

The paper has the following main contributions: it provides a general evaluation of intrusion-tolerant protocols in wireless networks; it proposes an algorithm for the speed agreement operations of a distributed cruise control, presenting the respective implementation and performance evaluation; it presents the first implementation and experimental evaluation of intrusion-tolerant protocols using mobile devices such as PDAs.

2. Distributed Cruise Control

Classical cruise control is a system designed to automatically control the speed of a vehicle. The driver sets a speed, and the system assumes automatic control of the throttle of the vehicle. The same speed is maintained despite changes in road conditions (e.g., inclinations).

Distributed cruise control is an extension of this functionality for platoons. A *platoon* is a tightly spaced vehicle group formation, and is aimed at improved highway efficiency (e.g., lane throughput, fuel consumption, etc.) and passenger comfort as much as possible. In the context of this work, a platoon is defined as a group of vehicles on a highway traveling at close distances from each other and heading to some common destination. This is the case, for instance, of platoons of heavy-duty freight trucks, which travel together from the departure to the arrival point [22].

The main difference between classical cruise control and distributed cruise control is that instead of having each vehicle defining its speed individually, the platoon members communicate in a wireless ad-hoc environment in order to agree on a common speed. Once chosen, this baseline speed is enforced by all vehicles, though not strictly. Vehicles are allowed to deviate from this baseline speed in order to maintain safe distances from each other.

This paper presents an algorithm that solves the *speed agreement* problem. It is designed as a service that can be used by a sophisticated distributed cruise control application to reach agreement on a baseline speed.

The speed agreement algorithm is arranged at the top of a protocol stack as depicted in Figure 1. The protocols be-




Figure 1. Protocol stack

low the speed agreement are used as primitives by this algorithm. At the bottom of the stack, there are the *reliable broadcast* and *echo broadcast* protocols. Reliable broadcast ensures that, upon a broadcast, all correct processes either deliver the same message or no message at all [3]. Echo broadcast is a weaker, but more efficient, version of reliable broadcast [18]. In the case where the sender is corrupt, it does not guarantee that all correct processes will deliver the message. It only ensures that the subset of correct processes that deliver, will do it for the same message.

Atop these broadcast primitives is binary consensus. This protocol, being the simplest form of consensus, is where randomization is applied to circumvent the FLP result. With it, processes can agree on a binary value. Two distinct randomized binary consensus protocols were implemented. The first, Bracha's binary consensus, is a local coin protocol that has a high time and communication complexity but avoids the use of expensive cryptography [3]. The second, ABBA binary consensus, is a shared coin protocol that has a low time and communication complexity but uses several asymmetric cryptographic primitives [4]. These protocols were subject to an experimental comparison in (wired) Ethernet LANs in a previous paper [15].

Above binary consensus is *multi-valued consensus*, which allows processes to propose and decide on a value with an arbitrary domain. Depending on the proposals, the decision is either one of the proposed values or a default value $\perp \notin \mathcal{V}$.

At the top of the stack is *vector consensus*. This protocol allows processes to agree on a vector with a subset of the proposed values. It ensures that every correct process decides on the same vector V of size n ; if a process p_i is correct, then the vector element $V[i]$ is either the value proposed by p_i or the default value \perp , and at least $f + 1$ elements of V were proposed by correct processes. The implemented protocol is the one described in [8].

3. System model

This section defines the system model for the distributed cruise control application. It applies to the speed agreement protocol and to all the protocols below in the stack.

A platoon is formed by a group of n vehicles (also called *processes* or *nodes*) $P = \{p_1, p_2, \dots, p_n\}$, where $n \geq 4$, and each pair of vehicles (p_i, p_j) shares a secret key s_{ij} . It is assumed that a platoon is formed at some departure point and travels together to some arrival point. Split and merge maneuvers are outside the scope of our problem and are not considered for simplicity.

Only vehicles belonging to the platoon can participate in the execution of the protocols. Nevertheless, some members of the platoon may be either ill-intentioned or corrupted by external entities. A vehicle is said to be *correct* if it follows the protocols until arrival. Otherwise, it is said to be *corrupt*. The only assumption on the behavior of corrupt vehicles/processes is that they follow the principle of no self-harm. This means that a vehicle cannot force its collision with other vehicles or take any action representing danger to its occupants, either by means of explicit car control (e.g., sudden breaking or acceleration) or transmission of incorrect information. They can, however, force collisions between other vehicles, or take any action they like to disrupt the correct operation of the platoon. There is a limit to the number corrupt vehicles that can exist in a platoon. In a platoon with n vehicles, at most $f = \lfloor \frac{n-1}{3} \rfloor$ of them can be corrupt.

Correct processes are assumed to be fully-connected by *authenticated reliable channels* that provide authenticity, reliability and integrity. Authenticity means that the recipient of a message knows who was the sender, reliability means that messages are eventually received, and integrity means that messages are delivered without modifications.

There are no assumptions about time, meaning that there are no bounds to the processing times or communication delays. Hence, the system is completely asynchronous.

4. Speed Agreement Algorithm

This section presents the speed agreement algorithm for intrusion-tolerant distributed cruise control. It allows vehicles inside a platoon to agree on a common baseline speed even if some of them attempt to thwart the correct platoon operation. The chosen baseline speed can be obtained using two different approaches, which are not necessarily mutually exclusive: a leader-based strategy, and a decentralized strategy.

Leader-based strategy. One of the platoon members, which is not necessarily the car at the head of the platoon, is defined a priori to be the leader, and disseminates the target speed to the rest of the platoon.

For this strategy, the challenge in terms of safety is to ensure that the leader reliably disseminates the speed value. The leader itself can be malicious, this means that mechanisms must be in place that keep the leader from disseminating contradictory speed values to different members of the platoon. It also must be ensured that all members of the

platoon eventually receive the target speed value and are not left in the dark since this could seriously affect convergence to a stable situation.

Therefore, one wants to guarantee that the information disseminated by the leader is delivered by all vehicles, it is not contradictory, and it is actually originated by the leader. A *reliable broadcast* protocol solves this problem since it ensures that every process receives the broadcasted value.

Decentralized strategy. Every member proposes a target speed, and an agreement protocol is executed to decide, based on or among the proposed values, a value to be used by all platoon members.

This second approach is potentially more immune to the scenario where some vehicle cannot guarantee the chosen speed. Since the decision is based upon the proposals from all members, it can, for instance, be the slowest car (under certain limits) to determine the target speed. It has to be ensured that the vehicles actually decide on the same speed, even if some of them try to disrupt this procedure by, for instance, disseminating contradictory proposals.

There are two protocols that in essence solve this problem: multi-valued consensus and vector consensus. *Multi-valued consensus* allows processes to decide on a value among the proposed ones. *Vector consensus* allows processes to agree on a vector composed by a subset of the proposed values. The latter is preferable since it provides more flexibility to the platooning application. It is preferable for the application to have access to a common vector composed by the proposals from the vehicles where it can apply some deterministic function to obtain to target speed (e.g., the lowest value, or the most common value), instead of being tied down to the single value returned by a multi-valued consensus protocol. For this reason, vector consensus is used by this strategy.

The algorithm. The speed agreement algorithm will apply both strategies complementarily. There will be a pre-defined leader – the car at the head of the platoon – that will set the target speed. In case some vehicle is unable to apply the stipulated speed, it can demand a decentralized agreement by having the vehicles execute a vector consensus. It is assumed that if, after a decentralized decision, a vehicle is unable to enforce the target speed, some kind of mechanism (e.g., manual intervention) ejects the car from the platoon. This procedure is, however, outside the scope of our application.

The speed agreement algorithm works by having the leader to reliably broadcast the new target speed. Then, all vehicles that can, assume the new target speed. If any vehicle cannot do this, it reliably broadcasts a message stating so. Upon receiving this message, all vehicles initiate a vector consensus. The new target speed is obtained by applying a deterministic function on the decision vector. All vehicles assume this new target speed, no exceptions.

5. Experimental Evaluation

This section evaluates the performance of the protocol stack of Figure 2 in 802.11 wireless LANs. First, the supporting protocols – reliable and echo broadcast through vector consensus – are evaluated independently via a set of micro-benchmarks. The experiments gauge the protocols in a number of environmental settings that are modeled as system parameters. This gives an insight into the impact of the various wireless environments on the performance of the protocols. Second, the speed agreement algorithm is evaluated under various types of faults that are injected into the stack to assess the resilience of the application in hostile environments.

Testbeds. The experiments were carried out on two different testbeds, *tb-emulab* and *tb-pda*.

The first, *tb-emulab*, was formed by 11 nodes from the Emulab network testbed [21]. Each node was a Pentium III PC with 600 MHz of clock speed and 256 MB of RAM, and contained a 802.11 a/b/g D-Link DWL-AG530 WLAN interface card, which was able to operate as an access point (AP). The operating system running on these nodes was Redhat Linux 9 with kernel version 2.3.34. The nodes were located on the same physical cluster and were, at most, a few meters distant from each other.

The second testbed, *tb-pda*, was formed by 7 HP hw6915 PDAs. These PDAs were equipped with an Intel PXA270 416 MHz processor, 64 MB of SDRAM, and integrated 802.11b WLAN. The operating system was Windows Mobile 5. The experiments were taken with the PDAs placed on the same table a few centimeters apart from each other. The access point used in this testbed was an Asus WL-320gE 802.11 b/g.

Implementation. The broadcast and consensus protocols were taken from the RITAS suite, which provides an implementation for Linux. These protocols were then ported to the Windows Mobile platform. The car platooning application was then developed for both platforms on top of the existing protocols. The reliable channels were implemented by the use of TCP for reliability, and the IPSec Authentication Header protocol for integrity.

Experimental Methodology. The performance metric utilized in the experiments was the *latency*. This metric is always relative to a particular process p_i . In the case of the consensus protocols, it is denoted as the interval of time between the moment p_i proposes a value to a consensus execution, and the moment p_i decides the consensus value. In the case of the broadcast protocols (and the speed agreement algorithm), it is the interval of time between the moment the sender (or leader) process, say p_i , initiates the execution of the protocol, and the moment p_i delivers the broadcasted value (or accepts a speed value).

The environmental system parameters are configurable parameters of the system that define specific execution environments. These are the *group size*, the *wireless standard and network bandwidth*, and the *network topology*. The *group size* defines the number of processes n in the system, and in our case it can take three values: 4, 7, and 10. The *wireless standard and network bandwidth* defines the amendment to the 802.11 WLAN standard used, and intrinsically defines the amount of bandwidth available in the network. Three WLAN standard are used: 802.11a, 802.11b, and 802.11g. The 802.11a and 802.11g standards provide 54 Mb/s bandwidth, and 802.11b provides 11 Mb/s. The *network topology* defines the way the network nodes communicate with each other. There are two types of network topology: ad-hoc and infrastructure. In the ad-hoc network topology the nodes communicate directly with each other with no access points. In the infrastructure network topology, all nodes communicate through an access point (AP). This kind of topology can also make sense in the context of car platooning since the vehicles may also communicate through a roadside infrastructural network.

The experiments were carried out the following way. A signaling machine, which does not participate in the execution of the protocols, is selected to conduct the experiment. It repeats the following procedure m times: it broadcasts a 1-byte UDP message to the n processes involved in the experiment. When a process receives one of these messages, it executes whatever protocol is relevant for the current experiment (the information about which protocol to execute is carried within the 1-byte UDP message). Processes record the latency value as described above, and send a 1-byte UDP message to the signaling machine indicating the termination of the execution of the protocol. The signaling machine, upon receiving n such messages, waits five seconds, and recommences the procedure. The *average latency* is obtained by taking the mean value of the sample of measured values.

5.1. Micro-benchmarks

The micro-benchmarks are a set of experiments that aim to access the impact of the several environmental settings on the performance of the protocols. This section presents four micro-benchmarks that evaluate specific settings. The first analyzes the impact of the wireless standard with different group sizes. The second, the effects of the computational capability of the processes. The third, the impact of the network topology. Finally, the fourth focuses on the binary consensus protocols. These protocols are interesting to look in more detail since they are the only ones that employ randomization, and do so using different strategies. In all the experiments, the message payload size of the broadcast and consensus protocols was set to 100 bytes. The only exception is binary consensus where 1-byte payloads are used (since the protocol only deals with binary values it does not

Wireless Standard	Group Size	Latency (ms)			
		Reliable Broadcast	Binary Consensus	Multi-val. Consensus	Vector Consensus
802.11b	$n = 4$	20.2	42.6	178.4	259.9
	$n = 7$	72.4	292.7	1581.7	2078.3
	$n = 10$	219.2	832.8	4678.9	6234.1
802.11g	$n = 4$	8.3	18.2	79.9	109.3
	$n = 7$	24.7	92.6	564.6	879.9
	$n = 10$	62.5	326.3	1608.1	2504.6
802.11a	$n = 4$	7.7	17.1	73.1	94.8
	$n = 7$	23.4	73.3	438.8	720.4
	$n = 10$	50.6	310.8	1340.5	1828.5

Table 1. Latency measurements for different wireless standards and group sizes in testbed *tb-emulab* in infrastructure mode.

make sense to have larger payloads). Except where is presented an explicit comparison between the two binary consensus protocols, the evaluated protocol is always Bracha’s binary consensus [3].

Wireless Standard and Group Size. This micro-benchmark evaluates the performance impact of both the wireless standard and the group size. The used testbed was *tb-emulab*. The network topology was set to infrastructure with one of the nodes acting exclusively as an access point. All possible wireless standard and group size settings were tested. The tested protocols were reliable broadcast, Bracha’s binary consensus, multi-valued consensus, and vector consensus.

Table 1 shows the obtained measurements. The relative cost of the protocols can be easily observed. It is completely congruent with their interdependencies within the stack. The greatest gap is from the binary consensus to the multi-valued consensus and is justified by the large messages the multi-valued consensus has to reliably broadcast to justify the proposed values [8]. The gaps from the reliable broadcast to the binary consensus, and from the multi-valued consensus to the vector consensus are smaller and directly related with the overhead incurred from the respective upper-layer protocols.

The performance impact of the wireless standard is mainly a consequence of the available bandwidth. The experiments with 802.11b (11 Mb/s) were significantly slower than the ones with 802.11g and 802.11a (54Mb/s). A reliable broadcast with four processes takes 8.3 ms on a 802.11g network, while on a 802.11b network this more than doubled to 20.2 ms. This pattern is roughly observed for all the experiments, while the difference becomes slightly more accentuated with larger group sizes.

Another interesting result is that the values obtained in the 802.11a experiments were consistently lower than the ones obtained in 802.11g, despite both standards being capable of achieving the same bandwidth. This difference is modest for the cheaper protocols and smaller group sizes. For instance, a reliable broadcast with four processes costs

Group Size	Testbed	Latency (ms)			
		Reliable Broadcast	Binary Consensus	Multi-valued Consensus	Vector Consensus
$n = 4$	<i>tb-emu.</i>	12	35	160	210
	<i>tb-pda</i>	26	211	320	374
$n = 7$	<i>tb-emu.</i>	36	154	972	1474
	<i>tb-pda</i>	52	1626	2555	3221

Table 2. Average latency in 802.11b ad-hoc network for both testbeds.

8.3 ms in 802.11g, and 7.7 ms in 802.11a, an almost negligible difference. However, as the protocols become more expensive and the group size increases (i.e., the network becomes more stressed), this difference becomes substantial. For vector consensus with ten processes, the cost is 2504.6 ms in 802.11g, and 1828.5 ms in 802.11a, which is a considerable difference.

Computational Capability. The second set of experiments measure how the computational capability of the individual nodes affects the performance of the protocols. Both testbeds were configured with the same system parameters: the wireless standard was set to 802.11b, the group size to 4 and 7 processes, and the network topology to ad-hoc mode. The computational capability in this contexts refers not just to the processing power of the CPU, but the whole local environment where the protocols are executed (e.g., hardware, drivers, operating system).

The obtained measurements are presented in Table 2. From the results it is clear that the computational characteristics of the individual nodes greatly affect the performance of the protocols. There is always a significant gap between the two testbeds for all the experiments. The average latency roughly doubles from *tb-emulab* to *tb-pda*, except for binary consensus where the difference is much greater.

In both testbeds it is observed that, at some point, a larger gap exists from one protocol to another. In *tb-pda* this happens from reliable broadcast to binary consensus, and in *tb-emulab* from binary consensus to multi-valued consensus. Being the system parameters equal for both testbeds, one must assume that it is the limited computational capability of the nodes in *tb-pda* that is responsible for the gap observed from reliable broadcast to binary consensus.

For testbed *tb-emulab*, the gap from binary consensus to multi-valued consensus is congruent with the messages the latter protocol has to exchange. It is interesting to note that the average latency measured in *tb-pda* gets progressively closer to that measured in *pda-emulab* as the cost of the protocols increases. What happens is that another degradation factor comes into play: the network bandwidth. So, as the communication cost of the protocols increases, the performance bottleneck shifts from the computational capability of the nodes to the network bandwidth. This is why the gap between the two testbeds tends to decrease as more stress is put on the network.

Group Size	Testbed	Latency (ms)			
		Reliable Broadcast	Binary Consensus	Multi-valued Consensus	Vector Consensus
$n = 4$	ad-hoc	26	211	320	374
	infra.	43	631	952	1190
$n = 7$	ad-hoc	52	1616	2555	3221
	infra.	138	7620	10062	12929

Table 3. Average latency for *tb-pda* with 4 and 7 processes.

Network Topology. This section looks at how the network topology of wireless networks impacts the performance of the protocols. For this experiment, measurements were taken in testbed *tb-pda* with 802.11b networks for both ad-hoc and infrastructure. The group size was set to 4 and 7 processes.

The measurements are presented in Table 3. The observation is that the operation in infrastructure mode does have a significant impact on performance. It introduces an additional delay into the communication between the processes since all data must be relayed through the AP.

The performance penalty in infrastructure mode remains essentially the same across all protocols despite their relative cost. Around 3 times with four processes, and 4 times with seven processes. For instance, in testbed *tb-pda*, a multi-valued consensus with four processes took 320 ms on average in ad-hoc mode, and 952 ms in infrastructure mode. With seven processes, it took 2555 ms in ad-hoc mode, and 10062 ms in infrastructure mode. So, a larger group also emphasizes a bit the degradation brought up by the presence of the AP.

These results demonstrate the high sensitivity these protocols have to the network latency, even more than the bandwidth, because of the large number of communication steps involved.

Binary Consensus Comparison. This section performs a more in-depth analysis of a key protocol in the stack: binary consensus. It compares the two different implementations of the protocol and presents some considerations about the performability of the two strategies employed by the protocols – one depends on the heavy use of public-key cryptography, and the other on abundant message exchanges. Testbed *tb-emulab* was used, and the following system parameters were tested: the group size for 4, 7, and 10 processes; the wireless standard for 802.11b, and 802.11g; and the network topology for ad-hoc, and infrastructure.

Given the features of the protocols it was expected for the Bracha protocol to outperform the ABBA protocol given more favorable network conditions, and to exist a certain point, as network conditions degrade, where the ABBA strategy would pay off to the point of being faster than the Bracha protocol. Table 4 presents the measurements observed for the various environmental settings tested.

The results confirm the expectations. The table shows

Group Size	Algorithm	Latency (ms)			
		802.11g		802.11b	
		ad-hoc	infra.	ad-hoc	infra.
$n = 4$	Bracha	11	18	35	43
	ABBA	147	148	146	160
$n = 7$	Bracha	43	94	154	293
	ABBA	210	211	211	270
$n = 10$	Bracha	95	326	415	833
	ABBA	290	311	301	717

Table 4. Average latency for binary consensus protocols in *tb-emulab*.

the measurements for the best network configuration (802.11g and ad-hoc), where the Bracha protocol is clearly faster than ABBA, even for a group size of ten processes – 95 ms against 290 ms, respectively. In remaining scenarios where the network conditions are not so good – either there is an AP or the standard is 802.11b, or both – there is a point from which ABBA outperforms Bracha. For the 802.11b/adhoc and 802.11g/infrastructure scenarios, this happens when the group size is ten, and for the 802.11b/infrastructure, which is the worst network configuration, this happens at $n = 7$. The conclusion is that Bracha is much faster with few processes and “good” network conditions, but it quickly degrades with the network capacity up to a point where the ABBA protocol, being more resilient to the network limitations, becomes faster.

5.2. Speed Agreement Algorithm

This section evaluates the speed agreement algorithm from Section 4. The application is evaluated in both testbeds and several scenarios are considered where the application is subject to various types of faults and different behaviors from the vehicles.

For testbed *tb-emulab*, two configurations are used, both with four nodes: one with a 802.11g ad-hoc network, and another with a 802.11g ad-hoc network. For *tb-pda*, the same number of nodes is used, but only on a 802.11b ad-hoc network. On both testbeds, the chosen binary consensus protocol is the Bracha’s protocol. The application is evaluated under three types of faultloads, and two distinct vehicle behaviors.

The three faultloads are normal, silent, and Byzantine. In the *normal* faultload, the vehicles follow the protocol correctly until termination. In the silent faultload, one of the vehicles does not communicate with the others. Finally, in the Byzantine faultload, f of the vehicles try their best to disrupt the execution of the platooning application. This is done at the binary consensus and multi-valued consensus protocols. At the binary consensus, they always propose zero, trying to force the processes to decide on this value. This has the effect of the multi-valued consensus to return a default value \perp , which in turn forces the vector consensus to execute another round. At the multi-valued consensus

layer, they always propose the default value, trying to force the processes to decide on the default value. The goal is also to force the vector consensus into another round.

The vehicle behavior is related to the way vehicles react to the speed agreement algorithm, and it can be either *cooperative* or *contentious*. In the cooperative behavior, all the vehicles accept the proposal of the leader. In the contentious behavior, one of the vehicles disagrees with the value proposed by the leader and forces a decentralized agreement.

Table 5 presents the results for all the possible combinations between faultload and vehicle behavior. The first observation is that the contentious behavior imposes a significant performance penalty. For instance, in testbed *tb-pda* (802.11b), an execution with cooperative behavior takes an average of 41.30 ms, and 715.45 ms with contentious behavior. This result is explained because the contentious behavior forces extra reliable broadcast and vector consensus executions. Even with this performance penalty, it is very interesting to note that the protocol can be executed quite fast in typical 802.11g network with moderately aged hardware. The executions with cooperative behavior in *tb-emulab* took around 3-4 ms in average, and around 60 ms with contentious behavior. These are promising results, specially considering that none of the protocols were optimized for wireless networks.

The observations about the wireless standard and the testbed impact are just confirmations of the results obtained in the previous sections. The testbed *tb-emulab* is faster than *tb-pda* under the same environmental conditions (802.11b and ad-hoc), and the 802.11g results are much better than the 802.11b.

The measurements about the faultload yield some interesting results. The *silent* faultload made the executions faster than the *normal* faultload. When one of the vehicles ceases communication, it alleviates the transmission medium, leaving more bandwidth available, and making the executions faster. Another interesting result is that the Byzantine failures did not have a noticeable impact on the performance of the algorithm. The resilience of the protocols extended to their performance, and confirms that the randomization techniques employed are not affected by Byzantine faults in terms of liveness, unlike most other approaches.

Feasibility of the Speed Agreement Algorithm. Despite the fact that many parameters would have to be taken into account to build a more complete control-based solution for a platooning application, we can reason about the possible impact of the communication delays affecting the behavior of the platoon. In particular, it is important to understand if the typical magnitude of protocol execution delays are adequate for this application.

To do that, let us consider a simple case, in which the platoon is stable and, at some point, the leader vehicle slows

Testbed/Standard	Faultload	Latency (ms)	
		Vehicle behavior	
		<i>Cooperative</i> (ms)	<i>Contentious</i> (ms)
<i>tb-pda</i> /802.11b	normal	41.30	715.45
	silent	29.24	501.11
	Byzantine	42.07	714.98
<i>tb-emulab</i> /802.11g	normal	4.21	64.31
	silent	3.81	57.09
	Byzantine	4.64	64.47
<i>tb-emulab</i> /802.11b	normal	32.71	579.54
	silent	24.80	407.72
	Byzantine	31.99	581.21

Table 5. Average latency for the speed agreement protocol (ad-hoc mode).

down, reducing its speed by 40 Km/h (we can assume this is done instantaneously). At this point, a new agreement would be started to propagate the new speed to all other vehicles. Clearly, there is a time bound after which the agreement result is useless, because the follower car would have approached too much and some safety mechanism would have been triggered to prevent a possible collision. Therefore, to ensure a smooth adaptation, the agreement protocol must terminate reasonably fast. But how fast?

For the mentioned speed reduction, the follower car will approach the leader at about 40 Km/h, that is, 11 m/s. Therefore, if some control algorithm to be executed by the vehicles is configured to stabilize with an inter-vehicle distance of about 15 meters, this would mean that a one second delay would be allowed for the agreement protocol execution, while keeping a safety distance margin of about 4 meters. This provides a rough idea of the values that would be in place when considering this platooning application, showing that our results are in the general case well within these margins. Consequently, it seems possible to consider intrusion-tolerant solutions in this context.

6. Conclusions

This paper presents a performance evaluation of randomized intrusion-tolerant agreement protocols in 802.11 wireless networks, and the application and evaluation of these protocols in a car platooning application scenario subject to various types of faultloads. This section summarizes some important results obtained in these evaluations.

The measurements taken in 802.11a/g networks (54 Mb/s) were considerably better than the ones taken in 802.11b (11 Mb/s), showing how the available bandwidth can affect the performance of the protocols.

The execution of the protocols is slightly but consistently faster in 802.11a against 802.11g. Despite both being capable of achieving the same maximum data rate, the typical data rate is higher in 802.11a networks.

The computational capability of the individual processes can represent a significant performance bottleneck. Nevertheless, as the cost of the protocols increase and more stress is put on the network, this bottleneck tends to shift from the

computational capability to the network bandwidth.

The introduction of an access point, and the consequential relay of all communication through it, imposes a general performance penalty on the protocols. The protocols are highly sensitive to the network delays.

Bracha's binary consensus is faster than ABBA binary consensus when the network conditions are better (i.e., higher bandwidth, lower latency). Nevertheless, there is a point, as network conditions degrade, at which ABBA outperforms Bracha's.

The car platooning application in particular shows that these protocols can be applied in practical settings. Given relatively old hardware (Pentium III PCs) and a typical 802.11g ad-hoc network, the speed agreement algorithm can be executed between four vehicles in a matter of a few milliseconds (3-4 ms) in normal conditions, and around 60 milliseconds if some car forces a decentralized agreement. Other interesting results to retain from the car platooning application pertain to the protocol behavior with faults. If some car does not communicate, the protocols actually execute faster. Additionally, the performance is unaffected even if some vehicle attempts to actively disrupt the execution of the protocols.

References

- [1] P. G. Argyroudis and D. O'Mahony. Secure routing for mobile ad hoc networks. *IEEE Communications Surveys*, 7(3):2–21, 2005.
- [2] F. Bonnet, P. Ezhilchelvan, and E. Vollset. Quiescent consensus in mobile ad-hoc networks using eventually storage-free broadcasts. In *Proceedings of the 21st Annual ACM Symposium on Applied Computing*, Apr. 2006.
- [3] G. Bracha. An asynchronous $\lfloor (n-1)/3 \rfloor$ -resilient consensus protocol. In *Proceedings of the 3rd ACM Symposium on Principles of Distributed Computing*, pages 154–162, Aug. 1984.
- [4] C. Cachin, K. Kursawe, and V. Shoup. Random oracles in Constantinople: Practical asynchronous Byzantine agreement using cryptography. In *Proceedings of the 19th ACM Symposium on Principles of Distributed Computing*, pages 123–132, July 2000.
- [5] M. Castro and B. Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems*, 20(4):398–461, Nov. 2002.
- [6] D. Cavin, Y. Sasson, and A. Schiper. Consensus with unknown participants or fundamental self-organization. In *Proceedings of the 3rd International Conference on Ad-Hoc Networks & Wireless*, pages 135–148, 2004.
- [7] G. Chockler, M. Demirbas, S. Gilbert, C. Newport, and T. Nolte. Consensus and collision detectors in wireless ad hoc networks. In *Proceedings of the 24th ACM Symposium on Principles of Distributed Computing*, July 2005.
- [8] M. Correia, N. F. Neves, and P. Veríssimo. From consensus to atomic broadcast: Time-free Byzantine-resistant protocols without signatures. *Computer Journal*, 41(1):82–96, Jan. 2006.
- [9] V. Drabkin, R. Friedman, and M. Segal. Efficient Byzantine broadcast in wireless ad-hoc networks. In *Proceedings of the IEEE International Conference on Dependable Systems and Networks*, June 2005.
- [10] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, Apr. 1985.
- [11] J. S. Fraga and D. Powell. A fault- and intrusion-tolerant file system. In *Proceedings of the 3rd International Conference on Computer Security*, pages 203–218, Aug. 1985.
- [12] F. Greve and S. Tixeuil. Knowledge connectivity vs. synchrony requirements for fault-tolerant agreement in unknown networks. In *Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 82–91, 2007.
- [13] Y.-C. Hu and A. Perrig. A survey of wireless ad hoc routing. *IEEE Security and Privacy*, 2(3):28–39, May/June 2004.
- [14] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. The SecureRing group communication system. *ACM Transactions on Information and System Security*, 4(4):371–406, Nov. 2001.
- [15] H. Moniz, N. F. Neves, M. Correia, and P. Veríssimo. Experimental comparison of local and shared coin randomized consensus protocols. In *Proceedings of the 25th IEEE Symposium on Reliable Distributed Systems*, pages 235–244, October 2006.
- [16] H. Moniz, N. F. Neves, M. Correia, and P. Veríssimo. Randomized intrusion-tolerant asynchronous services. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 568–577, June 2006.
- [17] M. K. Reiter. The Rampart toolkit for building high-integrity services. In *Theory and Practice in Distributed Systems*, volume 938 of *Lecture Notes in Computer Science*, pages 99–110. Springer-Verlag, 1995.
- [18] S. Toueg. Randomized Byzantine agreements. In *Proceedings of the 3rd ACM Symposium on Principles of Distributed Computing*, pages 163–178, Aug. 1984.
- [19] P. Veríssimo, N. F. Neves, and M. Correia. Intrusion-tolerant architectures: Concepts and design. In R. Lemos, C. Gacek, and A. Romanovsky, editors, *Architecting Dependable Systems*, volume 2677 of *Lecture Notes in Computer Science*, pages 3–36. Springer-Verlag, 2003.
- [20] E. Vollset and P. D. Ezhilchelvan. Design and performance-study of crash-tolerant protocols for broadcasting and reaching consensus in MANETs. In *Proceedings of the 24th IEEE Symposium on Reliable Distributed Systems*, pages 166–175, Oct. 2005.
- [21] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. In *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation*, pages 255–270. USENIX, Dec. 2002.
- [22] D. Yanakiev and I. Kanellakopoulos. Nonlinear spacing policies for automated heavy-duty vehicles. *IEEE Transactions on Vehicular Technology*, 47(4):1365–1377, Nov. 1998.