Vulnerabilities in Epidemic Forwarding

Alaeddine El Fawal

EPFL, I&C

1015 Lausanne, Switzerland

alaeddine.elfawal@epfl.ch

Jean-Yves Le Boudec

EPFL, I&C

1015 Lausanne, Switzerland

jean-yves.leboudec@epfl.ch

Kave Salamatian

EPFL, I&C

1015 Lausanne, Switzerland

kave.salamatian@epfl.ch

Abstract

We identify vulnerabilities in epidemic forwarding. We address broadcast applications over wireless ad-hoc networks. Epidemic forwarding employs several mechanisms such as inhibition and spread control, and each of them can be implemented using alternative methods. Thus, the existence of vulnerabilities is highly dependent on the methods used. We examine the links between them. We classify vulnerabilities into two categories: malicious and rational. We examine the effect of the attacks according to the number of attackers and the different network settings such as density, mobility and congestion. We show that malicious attacks are hard to achieve and their impacts are scenario dependent. In contrast, rational attackers always obtain a significant benefit. The evaluation is carried out using detailed realistic simulations over networks with up to 1000 nodes. We consider static scenarios, as well as vehicular networks.

1. Introduction

In this paper, we identify vulnerabilities in epidemic forwarding over ad-hoc networks. We are interested in broadcast applications such as chatting in a traffic jam or coupon advertisements [5]. The principle of epidemic forwarding is that nodes repeat with some probability the information they hear from others, thus propagating fresh information.

Epidemic forwarding employs several mechanisms such as inhibition, which prevents a node from forwarding over-sent or over-received packets in order to minimize redundancy. Each mechanism can be implemented using different alternative methods. Thus, the existence of vulnerabilities is highly dependent on the mechanisms employed and on the methods adopted to achieve them. We examine the links between these methods and the vulnerabilities.

We classify the vulnerabilities into two categories: malicious and rational. A malicious attacker harms other nodes and does not look for personal benefit. It aims at decreasing other nodes' throughput and/or spread. In contrast, a rational attacker aims at increasing its personal profit from the network. It tries to increase its throughput or save power.

We evaluate the vulnerabilities by simulations. We show that a malicious attacker does not have much effect in highly mobile networks, but it might be very harmful in static networks. In contrast, independently of mobility, the rational attacker increases dramatically its throughput. Attacks that are otherwise very harmful lose their efficiency in the presence of some epidemic forwarding mechanisms such as adaptive spread control and adaptive inhibition. Moreover, many elements such as the attacker position and the node density influence the malicious attacker's impact.

The simulations are carried out on networks with up to 1000 nodes using a JAVA implementation of the epidemic forwarding system in JIST-SWANS [1]. Beside static scenarios, we applied the epidemic forwarding system to the vehicular network using an extension of JIST-SWANS called STRAW [2], which provides a mobility model based on the operation of real vehicular traffic.

2. Epidemic Forwarding Mechanisms

In this section, we explain the different mechanisms used in epidemic forwarding, in order to understand their vulnerabilities.

2.1. Inhibition Mechanisms

Inhibition mechanisms aim at preventing nodes from forwarding over-sent or over-received packets in order to minimize redundancy. We classify them into two sets: rigid and adaptive. With the former set, the mechanisms cannot adapt themselves to different network settings: when the settings change, their parameters need to change. The adaptive mechanisms ensure a good performance in a wide range of settings without changing their parameters.

2.1.1. Rigid Inhibition. Within this set we find Gossip-based epidemic forwarding [7] where a node decides to forward a packet with a fixed probability p and drop it with $(1 - p)$. The value of p depends on the setting but Gossip does not involve any mechanism to adapt p .

2.1.2. Adaptive Inhibition. Within this set we distinguish between two methods.

2.1.2.a. Counter Based Inhibition: This method is essentially the one proposed in [8]. A packet stored in the

epidemic buffer has a counter called “Receive Count” incremented by 1 when a duplicate of this packet is received. Initially, i.e. when the packet is created by the application or received for the first time, the counter is set to 0. When the counter reaches a maximum value, the packet is discarded from the epidemic buffer. When a packet is transmitted, the value of Receive Count is lost.

2.1.2.b. Virtual Rate Based Inhibition: This method is proposed in [4]. With this method, a packet in the epidemic buffer is retransmitted with a probability that depends on its “virtual rate”; it is equal to $c_0 a^R b^S$ where c_0 is a constant (inverse of a time), R [resp. S] is the number of times this packet or a duplicate was received [resp. sent] and a and b are unit-less constants less than 1. Thus the virtual rate of a packet decreases exponentially with any send/receive event of the same packet. A scheduler decides which packet is selected next for transmission by the MAC layer; it serves packets with a rates not exceeding their virtual rates. Hence, a packet in the epidemic buffer, which has seen many send/receive events, is scheduled at a very low rate and it is more likely that it will be dropped by the used buffer management mechanism before being transmitted [4]. The constant c_0 is equal to the nominal packet rate of the MAC layer.

2.2. Spread Control Mechanisms

Spread control mechanisms are essential for epidemic forwarding, as the broadcast capacity does not scale with the population. Spread control can be implemented using one of the following methods.

2.2.1. Classic TTL. This is the method that comes by default with the Internet Protocol (IP). When a packet is created by a source and placed into the epidemic buffer, it receives a TTL value equal to some positive constant “MaxTTL”. When the packet is accepted for transmission by the MAC layer, the TTL field of the *transmitted* packet is equal to the value of the TTL field in the packet in the epidemic buffer, minus 1. The TTL field in the packet stored in the epidemic buffer is unchanged.

When a packet created by some other node is received for the first time at this node, the packet is delivered to the application, and the value of the TTL is screened. If it is equal to 0, it cannot be retransmitted and the packet is discarded. Else ($TTL \geq 1$), the packet is stored in the epidemic buffer, with TTL equal to the value present in the received packet. When and if the packet is later accepted for transmission by the MAC layer, the transmitted TTL field is equal to the stored TTL minus 1, and the stored TTL is unchanged.

2.2.2. Aging. This method is proposed in [4] in a different but essentially equivalent form. We give here a presentation that combines different options in one single framework. The method uses the TTL field like Classic TTL, but the TTL of a packet may be decremented while it is stored in the epidemic buffer, depending on receive and send events.

Formally, every packet in the epidemic buffer has an “age” field, which is a fixed decimal positive number less than 256. When a packet, created by some other node, is received by this node for the first time, its age is set to the complement to 255 of the received TTL: $\text{age} = 255 - \text{TTL}$. When a packet is transmitted, its stored age is incremented by a fixed amount K_0 and then its TTL is set to $255 - \text{age}$. When a duplicate packet is received, the received TTL is ignored but the stored age is incremented by K_1 : $\text{age} = \text{age} + K_1$. When *any* packet is received, the stored age of *all* packets in the epidemic buffer is incremented by K_2 : $\text{age} = \text{age} + K_2$. The node drops packets with age larger than 255.

2.3. Scheduler

Epidemic forwarding needs a scheduler for buffer management. To our knowledge, the only scheduler that is explicitly detailed in the literature is in [4]. It is used with the virtual-rate based inhibition (Sect. 2.1.2.b). It decides which packet in the epidemic buffer is selected for transmission, i.e. to be passed to the MAC layer. In order to ensure some level of fairness, the scheduler serves packets per source Id, using a processor sharing approach. Moreover, every packet should be served at a rate not exceeding its virtual rate computed in Sect. 2.1.2.b.

2.4. Control of Injection Rate

The only explicitly defined method to achieve control of injection rate is the one in [4]. It is used together with the aforementioned scheduler. The packets generated by the application at a given node are placed into the epidemic buffer, where they compete with the other packets for transmission (but with a larger virtual rate, having $R = S = 0$). The application rate is controlled by a windowing system : The number of outstanding packets the application is allowed to have in the epidemic buffer at this node is limited to -at most- 2 [4]; a packet is deleted from the epidemic buffer when a duplicate is received, which serves as implicit acknowledgment (Ack).

3. Attacks

In this section, we describe the vulnerabilities that are specific to epidemic forwarding. We distinguish between two types of attackers: malicious and rational. The former does not look for a personal benefit but aims to harm other nodes. In contrast, the latter seeks to increase its personal profit from the network. Most of the attacks are described by drawing (Figs. 1 and 2) using a generic example where the attacker is M and the victim in malicious case is A.

3.1. Malicious Attacks

A malicious attacker aims at decreasing the spread and/or the injection rate of the victim by exploiting vulnerabilities in epidemic forwarding mechanisms. In the following we identify five attacks and map them to their corresponding epidemic forwarding mechanisms.

M is the malicious node and *A* is the victim. We refer by *Self* to packets generated at the node transmitting them, by *Fwd* to packets forwarded by the node but generated by others and by *Fake* to packets that are generated by *M* but carrying the victim identity ($Id = A$). We will explain only the "Inhibit by forwarding and TTL" attack, as other attacks have similar explanation. In "Inhibit by Forwarding and TTL", *A* sends a *Self* packet *i* with $TTL = MaxTTL - 1$, that is received by *M* and *B*. *M* forwards the packet *i* (*Fwd: i*) 3 times with $TTL = 0$, the packet (*Fwd: i*) is received by *B* and *C*. Thus, the inhibition mechanism at *B* will inhibit packet *i*, as it is received 4 times, and *C* will drop the packet, as its $TTL = 0$.

Figure 1. Malicious attacks.

Figure 2. Sybil attack: *M* is the rational node.

3.1.1. Artificial High Density (AHD). In this attack, we exploit the adaptability of the spread control to the congest-

tion and node density. The attacker places itself close to the victim. It acts like any node: it has its self packets (packets that are generated at this node) to send and relays others packets. But it does not forward victim packets. By generating much traffic in the very close surrounding of the victim, the attacker incites the spread-control mechanism at the victim's good neighbors to react negatively and prevent the victim packets from going farther.

3.1.2. Inhibit by Forwarding (IbF) Attack. With IbF (Fig. 1), the attacker exploits the adaptive inhibition. It immediately forwards the victim packets a number of times (this number is called Attack-Persistency) to inhibit its neighborhood from forwarding the same packets (see Sect. 2.1.2). With the counter based inhibition (see Sect. 2.1.2.a), the Attack-Persistency is equal to the maximum value the counter can reach. With the virtual rate based inhibition, this Attack-Persistency should be large enough to make the corresponding virtual rate close to zero (in practice two times are enough).

3.1.3. Inhibit by TTL (IbTTL) Attack. This attack exploits the spread control using TTL. As the attacker receives a victim packet, it forwards it immediately with a $TTL = 0$. In Fig. 1, *B* and *M* receives a packet from *A* with $TTL = MaxTTL - 1$. *M* forwards it with $TTL = 0$ instead of ($MaxTTL - 2$). Hence, the attacker decreases the chance the packet has to travel beyond *C* as *B* is inhibited and *C* drops the packet. Even if *B* succeeds in forwarding the packet after *M*, this will change nothing with *C*. This example considers the use of "classic TTL" (see Sect. 2.2.1). With "aging", the attack is exactly the same. Note that in Fig. 1, *B* applies the first strategy (see Sect. 2.2.1) upon receiving a duplicate of the packet and thus it keeps the old TTL.

3.1.4. Inhibit by Forwarding and TTL (IbFTTL) Attack. This is a combination of IbF and IbTTL (Fig. 1). In this case *M* forwards the victim packets Attack-Persistency times with $TTL = 0$ to insure that the victim packets at *B* are well inhibited and thus the packet loses any chance of travelling beyond *C*.

3.1.5. Send on Behalf of the Victim (SoB) Attack. The attacker exploits the scheduler and the aging mechanism. In Fig. 1, *M* sends fake packets with *A*'s *Id*. As the scheduler serves packets per source *Id*, *A*'s packets are delayed in the epidemic buffer and they will be dropped either by the aging mechanism (they become too old) or by buffer overflow.

3.2. Rational Attacks

A rational attacker tries to increase its injection rate while maintaining large spread. In the following, we identify two rational attacks.

3.2.1. Do Not Cooperate (DNC) Attack. When a new packet is injected by the application at a given node, it is placed in the epidemic buffer, where it competes with pack-

ets received from other nodes. This competition prevents the application from injecting at the full rate allowed by the packet injection control mechanism because of the additional delay in the epidemic buffer. Thus, an attacker decides to not cooperate and to keep only its self packets (packets that are generated at this node) in the epidemic buffer. Note that, if the attacker tries to go beyond the allowed rate, its packets will be accumulated in other nodes, which are not able to serve them at the same rate. Thus, it risks killing its packets for the same reason as in Sect. 3.1.5.

3.2.2. Sybil Attack. We refer to the Sybil attacker as the node that forges multiple identities [3]. This is a well-known attack in networking, but the way it is exploited in this paper is new and very specific to epidemic forwarding. As the scheduler serves packets per source Id, the attacker sends its self packets with different Ids and thus it increases their share of the bandwidth. In Fig. 2, we present the scheduler as a process sharing approach where queues are per source Id. In this case, M's packets receive larger bandwidth share than A's packet at B.

4. Performance Evaluation

In this section, we evaluate the impact of the aforementioned attacks by simulation. We apply them in static scenarios, as well as in highly mobile networks. We consider vehicular mobility on the highway. Our metrics are based on the spread and the injection rate: a malicious attacker aims at reducing the victim spread and a rational one tries to increase its rate while maintaining large spread.

In our simulation we consider the epidemic forwarding system proposed in [4], called SLEF. To our knowledge, SLEF is the only complete system proposed for a wide range of settings. Furthermore, SLEF implements all epidemic forwarding mechanisms already discussed in Sect. 2: The virtual rate based inhibition, spread control by TTL and aging, injection rate control and the scheduler discussed in Sect. 2.3. The parameter values of the virtual rate based inhibition are $a = b = 0.15$, c_0 corresponds to 802.11b basic rate (1Mbps) with a packet length of 1500 bytes. As for the aging, we use $K_0 = K_1 = 25$ and $K_2 = 0.5$.

4.1. Settings

The victim is in the middle, the attackers start filling the positions around the victim according to their index in an increasing order: if one attacker, it fills position 1. If 2 attackers, they fill positions 1 and 2, and so on. R can be either 25m or 100m. $d = 25m$.

Figure 3. Malicious attackers positions.

With the static scenarios, we simulate from 200 up to 600 nodes uniformly distributed over a square of $500 \times 500 \text{ m}^2$, but, in most cases, we show the results only for 400 nodes as the others are similar. The transmission range is around 50 m (PDA transmission range).

In the case of a malicious attack, the victim is in the middle of the square and attackers take place around it as it is indicated in Fig. 3. We want to evaluate the impact of the distance between attackers and the victim. Therefore, the radius R in Fig. 3 can have one of two values: 25m and 100m. With the former, the attackers of the corresponding circle are within the transmission range of the victim and they are outside it with the latter.

In the case of a rational attack, there exists only one attacker, which is in the middle.

The network can be either congested, where all nodes are sources sending at full rate (capacity allowed by the channel) or non-congested, where the victim is the only source in the network and it is sending at full rate. Beside the victim, only attackers can act as sources in the non-congested scenario, based on the attack they want to achieve.

In the following we will use the following notations: “close” [resp. “far”] to indicate that R is equal to 25m [resp. 100m] and “one” [resp. “all”] to indicate that the network is non-congested [resp. congested].

As for the mobile scenario, we simulate 1000 vehicles in an urban two-lane road. The speed limit is 80 km/h. The car density is 12.5 cars/km in each direction. The transmission range is 300m, which is typical for vehicular network.

Our simulations are carried out through JIST-SWANS [1], an open source simulator for ad hoc networks. The MAC layer is a very accurate implementation of 802.11b in DCF mode with the basic rate of 1 Mbps as we transmit in broadcast (pseudo-broadcast [4]). As for the radio, we use the capture effect to approach the real WIFI cards that all implement it [6]. We consider fading channels with free space path-loss. As for the mobile network, we use an extension of JIST-SWANS called STRAW [2], which simulates the vehicular traffic and provides a mobility model based on the operation of real vehicular traffic.

4.2. Static Scenarios

4.2.1. Malicious Attacks.

4.2.1.a. AHD: The results are shown in Fig. 4. Let us begin with the “all” scenario (see Sect. 4.1) where the attackers are sources and act as any other node, except that they do not forward victim packets. In the “close” case, the impact of the attack is considerable in both scenarios, “all” and “one”, and it increases with the number of attackers. In contrast, in the “far” + “all” scenario, the attackers do not have a major impact; the reason is twofold: (1) the attackers are far from the victim and thus they do not increase the density as much as in the “close” scenario; (2) the inhibition mechanism is adapted. Indeed, the attackers are numerous and

Figure 4. Malicious attacks in static scenario.

Figure 5. Rate of a victim facing SoB attack.

they cooperate in forwarding all packets except the victims, hence they inhibit their neighbors from forwarding packets except those of the victim. Thus, the increase in the victim spread, which we notice for 8 attackers, is due to the fact that victim packets are less inhibited than others. If the inhibition were rigid, we expect that AHD would have more impact on the victim. In the “far” + “one” scenario, attackers are still injecting new packets in the network as before. They reduce considerably the victim spread.

4.2.1.b. IbF: The attackers do not generate fresh packets: their role is merely to forward victim packets as it is explained in Sect. 3.1.2. The results are shown in Fig. 4. It is clear that IbF does not achieve its goal. This can be explained as follows: When an attacker receives a new victim packet, it immediately forwards it Attack-Persistence times (if the MAC layer allows). If the same attacker receives another victim packet, before it finishes forwarding the previous packet, it cancels the previous and it starts anew with the newest. Thus, let us consider a scenario that happens frequently. The victim sends a new packet. The attacker forwards it immediately. The victim receives a duplicate of its self packet and considers it as an implicit Ack. Hence, it injects a new self packet that will be received by the attacker before finishing the forwarding process and it will be

received by other attackers even before beginning the forwarding process. Thus, all attackers cancel the previous packet, which explains why it is not inhibited.

4.2.1.c. IbTTL: Our implementation of IbTTL is similar to the one of IbF with the difference that it modifies the TTL before forwarding, as it is explained in Sect. 3.1.3. This attack is more harmful than IbF. The attacker needs to forward the packet only once with TTL = 0. Thus, nodes that receive the packet from the attacker for the first time are not able to forward it due to its TTL. This makes the difference with IbF, which needs to forward several times to inhibit the packet in its neighborhood.

4.2.1.d. IbFTTL: This attack has approximately the same impact as IbTTL, which is to be expected as IbF has little effect on the victim.

4.2.1.e. SoB: The attackers send only fake packets at full rate. Fig. 4 shows a significant decrease in spread and rate. The spread reduction is due to the fact that victim packets are killed in the epidemic buffers before being forwarded, which is due to the delay caused by the fake packets (for more explanation see Sect. 3.1.5). Moreover, the decrease in rate is due to the delay of the implicit Ack that controls the injection rate as explained in Sect. 2.4.

From what we have seen in this section, we can conclude that the attackers are not able to harm the victim in the presence of mobility for two reasons. The first is that the impact of the attackers is very position-dependent. The second is that, even with the most harmful attack, the attackers could reduce the spread of the victim, but its packets still reach a few tens of nodes. If these nodes are mobile, they will carry the victim packets beyond the barrier imposed by the attacker. This conclusion is well verified later in the vehicular network scenario.

Figure 6. Rational attacks in static scenario.

4.2.2. Rational Attacks.

4.2.2.a. DNC: We evaluate the impact of DNC only in the “all” scenario, where increasing the injection rate is a challenge. In the “one” scenario, the attacker is the only source in the network and he has the entire network capacity, thus it is meaningless to evaluate its impact in this case. In Fig. 6, the performance of a DNC attacker is compared with a well-behaved node that is very close to it and

thus they both experience the same network conditions. We show the spread of both nodes and the rate ratio (DNC over well-behaved). The DNC rate is four times larger than a well-behaved node. But, surprisingly, the DNC spread is much larger when the network is very dense (600 nodes). The reason is as follows: The attacker does not forward others' packets. Thus, when it receives others' packets, it drops them without updating the age of its self packets in the epidemic buffer. Hence, the age of its self packets does not increase during their stay in its epidemic buffer by K_2 (see Sect. 2.2.2), which allows them to travel farther.

4.2.2.b. Sybil: We evaluate the impact of Sybil in only "all" scenario for the same reason as with DNC. The attacker uses five different identities. In addition, it does not forward others' packets. So, our implementation is in fact a combination of both attacks, Sybil and DNC, explained in Sect. 3. This implementation gives the attacker a much larger advantage than using DNC alone (up to 10 times larger than a well-behaved node and 2.5 larger than the DNC attacker), which explains the impact of Sybil alone.

4.3. Vehicular Network Scenario

Figure 7. Vehicular Network. Malicious (a) and Rational (b) attacks, "all" scenario.

In this scenario, nodes are highly mobile and the position of the victim is not known. Thus, the attackers are chosen randomly. Beside the attackers, the network contains 1000 well-behaved nodes, all of them are sources ("all" scenario). All nodes cross the same urban road.

4.3.1. Malicious Attacks. Fig. 7(a) shows the impact of malicious attacks. The spread of the victim is drawn according to number of attackers. The Attack-Persistence of IbF and IbFTTL is 2. Other values give the same results. As we notice, the effect of the attackers is negligible even in the presence of 100 attackers. In the most harmful case, the IbFTTL attacker reduces the victim spread from 50 to 30 nodes, which is not significant. This can be explained by the presence of the spread control mechanism; the attacker can affect the victim only if their spreads interfere, i.e. there exist common nodes that receive the attacker and the victim packets. And the amount of harm is proportional to the amount of interference. As the spread is limited by

the spread control mechanism, this interference is not considerable and does not happen frequently.

4.3.2. Rational Attacks. Contrary to malicious attacks, rational attacks are still powerful even in highly mobile network. The results are shown in Fig. 7(b). Sybil still ensures higher gain than DNC.

5. Conclusions

We identify vulnerabilities that are specific to epidemic forwarding over wireless ad-hoc networks. We classify these vulnerabilities into two categories: malicious and rational. We evaluate their impact according to the number of attackers and the different network settings. We find that the impact of malicious attacks depends on the position of the attacker relative to the victim, the network density, the traffic load and mobility. In static scenarios, we identify the attacks that reduce dramatically the victim spread, whereas the harm of other attacks is reduced due to the adaptive inhibition and the injection rate control. In highly mobile vehicular network, the impact of malicious attacks are minimized due to the spread control.

We have studied the rational case in presence of only one attacker in the network. The attacker could achieve considerable profit in all scenarios.

Our work can be extended in different directions. We plan to examine the impact of the presence of several rational attackers on the network. Another extension is to find solutions to recover from these vulnerabilities.

References

- [1] Java in simulation time / scalable wireless ad hoc network simulator, jist/swans, <http://jist.ece.cornell.edu/>.
- [2] Street random waypoint / vehicular mobility model for network simulations, straw, <http://www.aqualab.cs.northwestern.edu/projects/straw/>.
- [3] J. R. Douceur. The sybil attack. In *The 1st International Workshop on Peer-to-Peer Systems (IPTPS'02)*, March 2002.
- [4] A. El Fawal, J.-Y. Le Boudec, and K. Salamatian. Self-Limiting Epidemic Forwarding. Technical Report LCA-REPORT-2006-126, EPFL, 2006.
- [5] A. Garyfalos and K. Almeroth. Coupons: Wide scale information distribution for wireless ad hoc networks. In *IEEE Global Telecommunications Conference (GLOBECOM) Global Internet and Next Generation Networks Symposium Dallas, Texas, USA*, pages 1655–1659, December 2004.
- [6] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b, berlin, germany. In *IEEE International Conference on Network Protocols (ICNP 04)*, pages 252–261, October 2004.
- [7] S.-D. Modiano, E. and G. Zussman. Maximizing throughput in wireless networks via gossiping. In *ACM SIGMETRICS / IFIP Performance'06*, June 2006.
- [8] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. In *Mobicom, Seattle, Washington, United States, August 15 - 19, 1999*, pages 151–162.