ISSN (Online) 2278-1021
ISSN (Print) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 4, Issue 9, September 2015

Vampire Attacks: Draining Life from Wireless Ad Hoc Sensor Networks

V.Kalaivani[1], M.Sumathi[2]

MCA, M.Phil, Mahendra Arts and Science College, Nammakal, Tamil Nadu[1]

M.Sc., M.Phil, Mahendra Arts and Science College, Nammakal, Tamil Nadu[2]

Abstract: Ad hoc low-power wireless networks are an exciting research direction in sensing and pervasive computing. Prior security work in this area has focused primarily on denial of communication at the routing or medium access control levels. This paper describes resource depletion attacks at the routing protocol layer, which permanently disable networks by quickly draining nodes' battery power. These "Vampire" attacks are not specific to any specific protocol, but rather rely on the properties of many popular classes of routing protocols. We find that all examined protocols are susceptible to Vampire attacks, which are devastating, difficult to detect, and are easy to carry out using as few as one malicious insider sending only protocol-compliant messages. In the worst case, a single Vampire can increase network-wide energy usage by a factor of O(N), where N in the number of network nodes. We discuss methods to mitigate these types of attacks, including a new proof-of-concept protocol that provably bounds the damage caused by Vampires during the packet forwarding phase.

Index Terms: Ad hoc networks, sensor networks, wireless networks.

1. INTRODUCTION

AD hoc wireless sensor networks (WSNs) promise exciting new applications in the near future, such as ubiquitous on-demand computing power, continuous connectivity, and instantly deployable communication for military and first responders. Such networks already monitor environmental conditions, factory performance, and troop deployment, to name a few applications. As WSNs become more and more crucial to the everyday functioning of people and organizations, availability faults become less tolerable lack of availability can make the difference between business as usual and lost productivity, power outages, environmental disasters, and even lost lives; thus high availability of these networks is a critical property, and should hold even under malicious conditions. Due to their ad hoc organization, wireless ad hoc networks are particularly vulnerable to denial of service (DoS) attacks, and a great deal of research has been done to enhance survivability.

In this paper, we consider how routing protocols, even those designed to be secure, lack protection from these attacks, which we call Vampire attacks, since they drain the life from networks nodes. These attacks are distinct from previously studied DoS, reduction of quality (RoQ), and routing infrastructure attacks as they do not disrupt immediate availability, but rather work over time to entirely disable a network. While some of the individual attacks are simple, and power draining and resource exhaustion attacks have been discussed before prior work has been mostly confined to other levels of the protocol stack, e.g., medium access control (MAC) or application layers, and to our knowledge there is little discussion, and no thorough analysis or mitigation, of routing-layer resource exhaustion attacks. Vampire attacks are not protocol-specific, in that they do not rely on design properties or implementation faults of particular routing protocols, but rather exploit general properties of protocol classes such as link-state, distance vector, source routing, and geographic and beacon routing. Neither do these attacks rely on flooding the network with large amounts of data, but rather try to transmit as little data as possible to achieve the largest energy drain, preventing a rate limiting solution. Since Vampires use protocol-compliant messages, these attacks are very difficult to detect and prevents. This paper makes three primary contributions. First, we thoroughly evaluate the vulnerabilities of existing protocols to routing layer battery depletion attacks. We observe that security measures to prevent Vampire attacks are orthogonal to those used to protect routing infrastructure, and so existing secure routing protocols such as Arianne, SAODV, and SEAD do not protect against Vampire attacks. Existing work on secure routing attempts to ensure that adversaries cannot cause path discovery to return an invalid network path, but Vampires do not disrupt or alter discovered paths, instead using existing valid network paths and protocol-compliant messages. Protocols that maximize power efficiency are also inappropriate, since they rely on cooperative node behavior and cannot optimize out malicious action. Second, we show simulation results quantifying the performance of several representative protocols in the presence of a single Vampire (insider adversary). Third, we modify an existing sensor network routing protocol to provably bound the damage from Vampire attacks during packet forwarding.

Definition

1.2 Overview

In this paper, we present a series of increasingly damaging Vampire attacks, evaluate the vulnerability of several

ISSN (Online) 2278-1021
ISSN (Print) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 4, Issue 9, September 2015

example protocols, and suggest how to improve resilience. In source routing protocols, we show how a malicious packet source can specify paths through the network which are far longer than optimal, wasting energy at intermediate nodes who forward the packet based on the included source route. In routing schemes, where forwarding decisions are made independently by each node (as opposed to specified by the source), we suggest how directional antenna and wormhole attacks can be used to deliver packets to multiple remote network positions, forcing packet processing at nodes that would not normally receive that packet at all, and thus increasing network-wide energy expenditure. Lastly, we show how an adversary can target not only packet forwarding but also route and topology discovery phases if discovery messages are flooded, an adversary can, for the cost of a single packet, consume energy at every node in the network.

In our first attack, an adversary composes packets with purposely introduced routing loops. We call it the carousel attack, since it sends packets in circles as shown in Fig. 1a. It targets source routing protocols by exploiting the limited verification of message headers at forwarding nodes, allowing a single packet to repeatedly traverse the same set of Nodes.

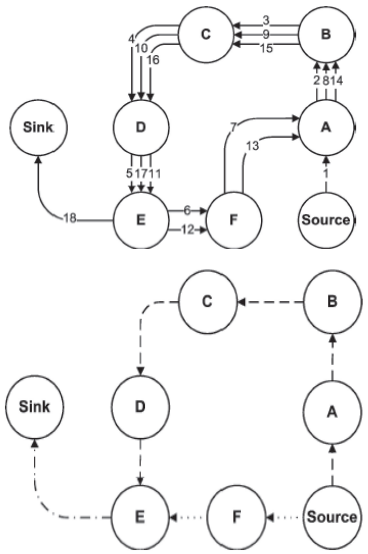

Fig.1.Malicious route construction attacks on source routing: carousel attack (a) and stretch attack (b).

We call this the stretch attack, since it increases packet path lengths, causing packets to be processed by a number of nodes that is independent of hop count along the shortest path between the adversary and packet destination. Results show that in a randomly generated topology, a single attacker can use a carousel attack to increase energy consumption by as much as a factor of 4, while stretch attacks increase energy usage by up to an order of magnitude, depending on the position of the malicious node. The impact of these attacks can be further increased by combining them, increasing the number of adversarial nodes in the network, or simply sending more packets. Although in Networks that do not employ authentication or only use end-to-end authentication, adversaries are free to replace routes in any overheard packets, we assume that only messages originated by adversaries may have maliciously composed routes.

We explore numerous mitigation methods to bound the damage from Vampire attacks, and find that while the carousel attack is simple to prevent with negligible overhead, the stretch attack is far more challenging. The first protection mechanism we consider is loose source routing, where any forwarding node can reroute the packet if it knows a shorter path to the destination. Unfortunately, this proves to be less efficient than simply keeping global network state at each node, defeating the purpose of source routing. In our second attempt, we modify the protocol from to guarantee that a packet makes progress through the network. We call this the no-backtracking property, since it holds if and only if a packet is moving strictly closer to its destination with every hop, and it mitigates all mentioned Vampire attacks with the exception of malicious flooded discovery, which is significantly harder to detect or prevent. We propose a limited topology discovery period ("the night," since this is when vampires are most dangerous), followed by a long packet forwarding period during which adversarial success is provably bounded. We also sketch how to further modify the protocol to detect Vampires during topology discovery and evict them after the network converges (at "dawn").

2. ATTACKS USING STATELESS PROTOCOLS

Here, we present simple but previously neglected attacks on source routing protocols, such as DSR . In these systems, the source node specifies the entire route to a destination within the packet header, so intermediaries do not make independent forwarding decisions, relying rather on a route specified by the source. To forward a message, the intermediate node finds itself in the route (specified in the packet header) and transmits the message to the next hop. The burden is on the source to ensure that the route is valid at the time of sending, and that every node in the route is a physical neighbor of the previous route hop. This approach has the advantage of requiring very little forwarding logic at intermediate nodes, and allows for entire routes to be sender authenticated using digital signatures, as in Ariadne.

We evaluated both the carousel and stretch attack in a randomly generated 30-node topology and a single randomly selected malicious DSR agent, using the ns-2 network simulator [1]. Energy usage is measured for the minimum number of packets required to deliver a single message, so sending more messages increases the strength of the attack linearly until bandwidth saturation.

We independently computed resource utilization of honest and malicious nodes and found that malicious nodes did not use a disproportionate amount of energy in carrying out the attack. In other words, malicious nodes are not driving down the cumulative energy of the network purely by their own use of energy. Nevertheless, malicious node energy consumption data are omitted for clarity[3]. The attacks are carried out by a randomly selected adversary

ISSN (Online) 2278-1021
ISSN (Print) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 4, Issue 9, September 2015

using the least intelligent attack strategy to obtain average expected damage estimates.

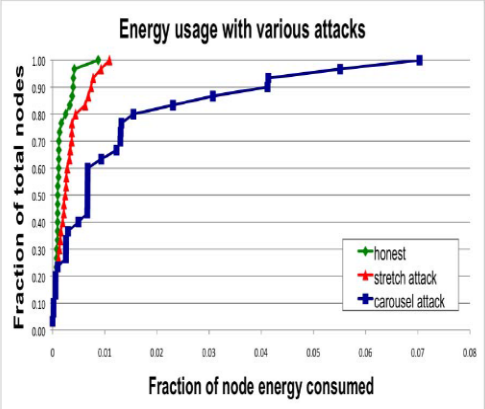

Fig. 2. Node energy distribution under various attack scenarios.

Per-node energy usage under both attacks is shown in Fig. 2. As expected, the carousel attack causes excessive energy usage for a few nodes, since only nodes along a shorter path are affected. In contrast, the stretch attack shows more uniform energy consumption for all nodes in the network, since it lengthens the route, causing more nodes to process the packet. While both attacks significantly network-wide energy usage, individual nodes are also noticeably affected, with some losing almost 10 percent of their total energy reserve per message. Fig. 3a diagrams the energy usage when node 0 sends a single packet to node 19 in an example network topology with only honest nodes. Black arrows denote the path of the packet. Carousel attack. In this attack, an adversary sends a packet with a route composed as a series of loops, such that the same node appears in the route many times. This strategy can be used to increase the route length beyond the number of nodes in the network, only limited by the number of allowed entries in the source route.2 An example of this type of route is in Fig. 1a. In Fig. 3b, malicious node 0 carriesout a carousel attack, sending a single message to node 19 (which does not have to be malicious). Note the drastic increase in energy usage along the original path.

2.1. Existing system

Existing work on secure routing attempts to ensure that adversaries cannot cause path discovery to return an invalid network path, but Vampires do not disrupt or alter discovered paths, instead using existing valid network paths and protocol compliant messages. Protocols that maximize power efficiency are also inappropriate, since they rely on cooperative node behavior and cannot optimize out malicious action.

2.2. Disadvantages of existing system

- ✓ Power outages
- ✓ Due to Environmental disasters, loss in the information
- ✓ Lost productivity
- ✓ Various DOS attacks
- ✓ Secure level is low
- ✓ They do not address attacks that affect long-term availability.

3. ATTACKS USING STATEFUL PROTOCOLS

We now move on to stateful routing protocols, where network nodes are aware of the network topology and its state, and make local forwarding decisions based on that stored state. Two important classes of stateful protocols are link-state and distance-vector. In link-state protocols, such as OLSR [2], nodes keep a record of the up-or-down state of links in the network, and flood routing updates every time a link goes down or a new link is enabled. Distancevector protocols like DSDV keep track of the next hopto every destination, indexed by a route cost metric, e.g., the number of hops. In this scheme, only routing updates that change the cost of a given route need to be propagated. Routes in link-state and distance-vector networks are built dynamically from many independent forwarding decisions, so adversaries have limited power to affect packet forwarding, making these protocols immune to carousel and stretch attacks. In fact, any time adversaries cannot specify the full path, the potential for Vampire attack is reduced. However, malicious nodes can still misforward packets, forcing packet forwarding by nodes who would ot normally be along packet paths. For instance, an adversary can forward packets either back toward the source if the adversary is an intermediary, or to a nonoptimal next hop if the adversary is either an intermediary or the source. While this may seem benign in a dense obstacle-free topology, worst case bounds are no better than in the case of the stretch attack on DSR. For instance, consider the special case of a ring topology: forwarding a packet in the reverse direction causes it to traverse every node in the network (or at least a significant number, assuming the malicious node is not the packet source but rather a forwarder), increasing our network wide energy consumption by a factor of O(N). While ring topologies are extremely unlikely to occur in practice, they do help us reason about worst case outcomes. Loose source routing performance compared to optimal, in a network with diameter slightly above 10. The dashed trend line represents expected path length when nodes store logN local state,and the solid trend line shows actual observed performance. Packet time to live (TTL) also limits route length, but it is set by the malicious sender. Intermediate nodes may able to reset it to a “reasonable” value, but it is unclear how to discover that value.

3.1.Stretch attack

Another attack in the same vein is the stretch attack, where a malicious node constructs artificially long source routes, causing packets to traverse a larger than optimal number of nodes. The network is composed of 30 nodes and a single randomly positioned Vampire. Results shown are based on a single packet sent by the attacker[4.

3.2 Directional antenna attack.

Vampires have little control over packet progress when forwarding decisions are made independently by each node, but they can still waste energy by restarting a packet in various parts of the network. Using a directional antenna adversaries can deposit a packet in arbitrary parts of the network, while also forwarding the packet locally. This consumes the energy of nodes that would not have

ISSN (Online) 2278-1021
ISSN (Print) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 4, Issue 9, September 2015

had to process the original packet, with the expected additional honest energy expenditure of O(N), where d is the network diameter, making d expected length of the path to an arbitrary destination from the furthest point in the network. This attack can be considered a half-wormhole attack , since a directional antenna constitutes a private communication channel, but the node on the other end is not necessarily malicious.7 It can be performed more than once, depositing the packet at various distan points in the network, at the additional cost to the adversary for each use of the directional antenna. Packet Leashes cannot prevent this attack since they are not meant to protect against malicious message sources, only intermediaries[5].

3.3. Malicious discovery attack

Another attack on all previously mentioned routing protocols (including stateful and stateless) is spurious route discovery. In most protocols, every node will forward route discovery packets (and sometimes route responses as well), meaning it is possible to initiate a flood by sending a single message. Systems that perform as-needed route discovery, such as AODV and DSR, are particularly vulnerable, since nodes may legitimately initiate discovery at any time, not just during a topology change. A malicious node has a number of ways to induce a perceived topology change: it may simply falsely claim that a link is down, or claim a new link to a nonexistent node. Security measures, such as those proposed by Raffo et al. may be sufficient to alleviate this particular problem. Further, two cooperating malicious nodes may claim the link between them is down. However,nearby nodes might be able to monitor communication to detect link failure (using some kind of neighborhood update scheme). Still, short route failures can be safely ignored in networks of sufficient density. More serious attacks become possible when nodes claim that a longdistance route has changed.

This attack is trivial in open networks with unauthenticated routes, since a single node can emulate multiple nodes in neighbor relationships [16], or falsely claim nodes as neighbors. Therefore, let us assume closed (Sybil-resistant) networks where link states are authenticated, similar to route authentication in Ariadne or path-vector signatures in the reference. Now our adversary must present an actually changed route in order to execute the attack. To do this, two cooperating adversaries communicating through a wormhole could repeatedly announce and withdraw routes that use this wormhole, causing a theoretical energy usage increase of a factor of O(N) per packet. Adding more malicious nodes to the mix increases the number of possible route announce/withdrawal pairs. Packet Leashes cannot prevent this attack, with the reasoning being similar to the directional antenna attack—since the originators are themselves malicious, they would forward messages through the wormhole, and return only seemingly valid (and functional) routes in response to discovery. This problem is similar to route flapping in BGP, but while Internet paths are relatively stable, paths change frequently in wireless ad hoc networks, where nodes may move in and out of each other'srange, or suffer intermittent environmental effects. Since there may be no stable routes in WSNs (hence the need for ad hoc protocols), this solution would not be applicable.

4. COORDINATE AND BEACON-BASED PROTOCOLS

Some recent routing research has moved in the direction of coordinate- and beacon-based routing, such as GPSR and BVR which use physical coordinates or beacon distances for routing, respectively. In GPSR, a packet may encounter a dead end, which is a localized space of minimal physical distance to the target, but without the target actually being reachable (e.g., the target is separated by a wall or obstruction). The packet must then be diverted (in GPSR, it follows the contour of the barrier that prevents it from reaching the target) until a path to the target is available. In BVR, packets are routed toward the beacon closest to the target node, and then move away from the beacon to reach the target. Each node makes independent forwarding decisions, and thus a Vampire is limited in the distance it can divert the packet. These protocols also fall victim to directional antenna attacks in the same way as link-state and distance-vector protocols above, leading to energy usage increase factor of O(N) per message, where d is the network diameter. Moreover, GPSR does not take path length into account when routing around local obstructions, and so malicious misrouting may cause up to a factor of O(N) energy loss, where c is the circumference of the obstruction, in hops.

5. CLEAN-SLATE SENSOR NETWORK ROUTING

In this section, we show that a clean-slate secure sensor network routing protocol by Parno et al. ("PLGP" from here on) can be modified to provably resist Vampire attacks during the packet forwarding phase. The original version ofthe protocol, although designed for security, is vulnerable to Vampire attacks. PLGP consists of a topology discovery phase, followed by a packet forwarding phase, with the former optionally repeated on a fixed schedule to ensure that topology information stays current. (There is no ondemanddiscovery.) Discovery deterministically organizes nodes into a tree that will later be used as an addressing scheme. When discovery begins, each node has a limited view of the network—the node knows only itself. Nodes discover their neighbors using local broadcast, and form ever expanding "neighborhoods," stopping when the entire network is a single group. Throughout this process, nodes build a tree of neighbor relationships and group membership that will later be used for addressing and routing.. The attack is not very effective when using virtual wormholes (encrypted connections), since adversaries sending packets to each other would accomplish the same goal.

6. PROPOSED SYSTEM

This paper makes three primary contributions. First, we thoroughly evaluate the vulnerabilities of existing protocols to routing layer battery depletion attacks. We observe that security measures to prevent Vampire attacks

ISSN (Online) 2278-1021
ISSN (Print) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 4, Issue 9, September 2015

are orthogonal to those used to protect routing infrastructure, and so existing secure routing protocols such as Ariadne, SAODV and SEAD do not protect against Vampire attacks. Existing work on secure routing attempts to ensure that adversaries cannot cause path discovery to return an invalid network path, but Vampires do not disrupt or alter discovered paths, instead using existing valid network paths and protocol-compliant messages. Protocols that maximize power efficiency are also inappropriate, since they rely on cooperative node behavior and cannot optimize out malicious action. Second, we show simulation results quantifying the performance of several representative protocols in the presence of a single Vampire (insider adversary). Third, we modify an existing sensor network routing protocol to provably bound the damage from Vampire attacks during packet forwarding.

In proposed system we show simulation results quantifying the performance of several representative protocols in the presence of a single Vampire. Then, we modify an existing sensor network routing protocol to provably bound the damage from Vampire attacks during packet forwarding.

6.1. Advantages of proposed system:

- ✓ Protect from the vampire attacks
- ✓ Secure level is high
- ✓ Boost up the Battery power

7. CONCLUSION

In this paper, we defined Vampire attacks, a new class of resource consumption attacks that use routing protocols to permanently disable ad hoc wireless sensor networks by depleting nodes' battery power. These attacks do not depend on particular protocols or implementations, but rather expose vulnerabilities in a number of popular protocol classes. We showed a number of proof-of-concept attacks against representative examples of existing routing protocols using a small number of weak adversaries, and measured their attack success on a randomly generated topology of 30 nodes. Simulation results show that depending on the location of the adversary, network energy expenditure during the forwarding phase increases from between 50 to 1,000 percent. Theoretical worst case energy usage can increase by as much as a factor ofO(N) per adversary per packet, where N is the network size. We proposed defenses against some of the forwarding-phase attacks and described PLGPa, the first sensor network routing protocol that provably bounds damage from Vampire attacks by verifying that packets consistently make progress toward their destinations.

REFERENCES

[1] "The Network Simulator - ns-2," http://www.isi.edu/nsnam/ns,2012.
[2] I. Aad, J.-P. Hubaux, and E.W. Knightly, "Denial of ServiceResilience in Ad Hoc Networks," Proc. ACM MobiCom, 2004.
[3] G. Acs, L. Buttyan, and I. Vajda, "Provably Secure On-Demand Source Routing in Mobile Ad Hoc Networks," .
[4] T. Aura, "Dos-Resistant Authentication with Client Puzzles," Proc. Int'l Workshop Security Protocols, 2001.
[5] Eugene Y. Vasserman and Nicholas Hopper "Vampire Attacks: Draining Life from Wireless Ad Hoc Sensor Networks"- IEEE transactions on mobile computing, vol. 12, no. 2, February 2013.