

INTEGRATING INTRUSION DETECTION AND FAULT LOCALIZATION IN MANETS

Dan Sterne, David Balenson
SPARTA, Inc.
{dan.sterne,david.balenson}@sparta.com

Simon Tsang, Petros Mouchtaris
Telcordia Technologies, Inc.
{stsang,pmouchta}@telcordia.com

Maitreya Natu, Adarshpal Sethi
University of Delaware
{natu,sethi}@cis.udel.edu

ABSTRACT*

In this exploratory paper, we propose that intrusion detection and fault localization techniques in MANET environments (which are commonly separate systems) should work cooperatively. We argue that an integrated approach will exhibit improved accuracy, and also minimize system overheads and redundancy. Using detection of in-band wormhole attacks as an illustrative example, we outline how an integrated approach can better distinguish malicious network attacks from “normal” network delays and outages.

Keywords: integrated intrusion detection and fault localization, MANET, dynamic hierarchy, mobile ad hoc network, network security, cyber security, network management, wormhole.

1. MOTIVATION

Mobile ad hoc network (MANET) technology is a key enabler for the Army’s vision of Network Centric Warfare (NCW) [GA04]. MANETs have no fixed or static infrastructure and dynamically change their topology to respond to node mobility, RF obstructions, and changing application and mission needs. As a consequence of this paradigm shift, the usefulness of existing techniques for intrusion detection and fault localization in wired networks is limited. In prior work, we have described separate techniques for intrusion detection [SBC05] and fault localization [SS04a, SS04b, SS04c]. While examining the relative strengths and weaknesses of these approaches, we have concluded that each could benefit by incorporating aspects of the approaches and algorithms developed by the other.

A key motivation for integrating intrusion detection and fault localization is that some of the symptoms these subsystems analyze overlap and cannot be analyzed by just one or the other (see Figure 1). Instead, such

symptoms must be analyzed as potential indicators of both faults and intrusions. For example, congestion in a particular neighborhood could be a symptom of 1) a benign fault such as loss of radio connectivity or inadequate balancing of bandwidth-intensive clients and servers; or 2) an intrusion, such as a distributed denial of service (i.e., flooding) attack or a routing attack creates non-functional routes towards which application traffic is futilely directed.




Figure 1: Network symptoms requiring cooperative fault localization and intrusion detection

Since the faults and intrusions typically require different kinds of responses, the results of symptom analysis by these two subsystems should be compared and integrated. In some cases, the analysis itself may need to be integrated. For example, if it can be determined that the congestion is caused by excessive traffic rather than poor connectivity, then determining whether the cause is a fault or an intrusion may hinge on whether the traffic appears to be legitimate or fabricated. This concept is illustrated in Figure 2. Note that response techniques are beyond the scope of this paper.




Figure 2: Integrated intrusion detection and fault localization concept

* Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

We further suggest that an integrated approach will minimize network overhead and redundancy due to efficiencies in an integrated hierarchical structure.

While we are in the process of validating the integrated approach and its benefits, this paper outlines our vision. Section 2 summarizes challenges faced by intrusion detection and fault localization systems in MANET environments. Sections 3 and 4 briefly review our intrusion detection and fault localization techniques for MANETs, utilizing detection of *in-band wormhole attacks* [KS05] as an illustrative example. Section 5 discusses in more detail our proposed approaches for and benefits gained from integrating intrusion detection and fault localization techniques.

2. CHALLENGES IN MANETS

Intrusion detection and fault localization for MANETs is particularly challenging and existing tools are of limited use due to the following MANET characteristics:

- Lack of centralized infrastructure: In MANETs, there is no central point for performing analysis of collected data. Analysis needs to be performed to place nodes such that overhead traffic is minimized. Furthermore, selection of those nodes cannot be static because connectivity of nodes changes substantially with time.
- Dynamic nature of ad hoc networks: MANET networks are highly dynamic with nodes and network infrastructure (routing, security, configuration) that may move, be destroyed, or lose connectivity. “Normal” MANET operation will include short periods of node isolation, link breakage and other behaviors which traditional intrusion detection or fault localization techniques may report as due to hostile or unintentional attacks. MANET specific techniques must be able to distinguish these behaviors. Furthermore, detection nodes may take different roles depending on the network connectivity at the time. For example, at some points in time, a node may simply collect data, while at other times, if sufficiently connected to other nodes, it may play a more central role in collecting and analyzing such data.
- Highly constrained bandwidth and network resources: Intrusion detection and fault localization techniques must minimize the amount of overhead (network packets) they produce. Network traffic information must be processed locally rather than sent to a central point. To improve the effectiveness of detection and localization, the techniques should

cooperatively share processed information about observed network behaviors.

- Lack of traffic chokepoints: In MANETs, there are no natural traffic concentration points/chokepoints to observe traffic. Traffic inspection and analysis techniques need be fully distributed and a large number of nodes need to participate and cooperate in the detection and localization process.

3. INTRUSION DETECTION TECHNIQUE

In this section, we summarize some of our recent research on intrusion detection for MANETs, beginning with a discussion of our distributed architecture for cooperative detection [SBC05]. We then describe our prototype detector for in-band wormhole attacks in OLSR networks and explain how it leverages the architecture.

3.1 Cooperative Detection Architecture

The choice of an organizational model is fundamental to the architecture of any distributed system including systems in which distributed detection and correlation components must exchange observations to detect and localize certain kinds of attacks. The model we are investigating in our research on intrusion detection for MANETs is the *dynamic hierarchy* [SBC05]. The major advantage of a hierarchy is its potential scalability to large networks, since it can provide rapid and communications-efficient detection for local cooperative attack recognition, while still allowing data sharing for more widely-distributed cooperative intrusion detection algorithms. Unlike peer-to-peer (P2P) networks where communications overhead can rise by the square of the number of nodes, a hierarchical approach allows higher-layer nodes to selectively aggregate and reduce intrusion detection data as it is reported upward from the leaf nodes to a root. Moreover, a hierarchy can naturally align with the authority structure or chain-of-command that is common to many human organizations. In the dynamic hierarchy, this structure is represented by the upward flow of data to authoritative nodes at or near the root of the hierarchy, which dispatch directives down to lower levels.

In a battlefield environment, mobility and other factors will cause a MANET's topology to change continually, such that an initially-defined static hierarchy will soon be inefficient. Since both nodes and links will appear and disappear rapidly, a dynamic, topology-based hierarchy must adapt on an on-going basis. Nodes will communicate intrusion detection information most often with other nodes that are their parents or children in the

appear as if it originated from a nearby node. Similarly, all subsequent OLSR control messages and other captured data packets are forwarded through the wormhole tunnel and rebroadcast. As a result of these exchanges and subsequent propagation of link state via the OLSR protocol, all benign nodes in network soon perceive a link between the nodes 189 and 186 that does not truly exist. We call this form, the *extended* in-band wormhole because the endpoints of the wormhole link 189 and 186 *extend beyond the tunnel endpoints* 180 and 183.

The *self-contained in-band wormhole* is similar. The attackers forward their own HELLOs to each other, or they simply falsely report each other as 1-hop neighbors in their own HELLO messages. All other OLSR control messages and data packets arriving at each tunnel endpoint are forwarded through the tunnel to the opposite endpoint and are handled there as if they had been received from a true one-hop neighbor. The result is that other nodes in the network perceive a link between nodes 180 and 183 that does not exist. In the *self-contained* in-band wormhole, both the *wormhole endpoints* and the *tunnel endpoints* are located at nodes 180 and 183.

3.3 Approach to Detection

MANETS depend on wireless (radio) links, which are inherently noisy. To meet military communications security requirements, MANETs deployed on the battlefield are likely to have more limited bandwidth and range than civilian wireless networks. As a result, packets traveling even relatively short distances will need to be forwarded through multiple, low-capacity hops and will incur significant loss and delays.

Our approach to detecting in-band wormholes is based on the hypothesis that because of the cumulative effect of loss and delay, it will be difficult for attackers to make a multiple-hop tunnel behave exactly like a true single-hop link. Even if the attackers can easily trick the routing layer into believing that wormhole link is real, it will be difficult make the illusion completely convincing in other respects. Given this hypothesis, our approach to detecting wormholes is to measure end-to-end loss and delay for various paths in the network and determine whether these measurements are within the range expected for the path lengths reported via the routing protocol. Significant discrepancies may indicate malicious manipulation of the routing layer and the presence of a wormhole.

Detecting the presence of a wormhole may be possible by analyzing loss and delay measurements, e.g. between nodes 189 and 186 for extended and between 180 and

183 for self-contained wormholes. If nodes are controlled by an adversary, as they are for the self-contained case, we cannot rely on them to report their own misbehavior to an intrusion detection system. Instead, we will attempt to collect end-to-end measurements for longer paths that encompass the wormhole link, but are likely to have pairs of benign endpoints from which we can expect reliable reporting.

The drawback of using end-to-end measurements of longer paths is that when the measurements appear anomalous, we require additional information, to determine which hop in the path is the culprit. To mitigate this uncertainty, our approach is to measure only 3-hop paths, as this minimizes the end-to-end path length required to encompass a self-contained wormhole while providing two benign endpoints. We then compare sets of anomalous 3-hop paths for overlap and hence can pinpoint the attacking nodes. Anomalous paths are forwarded upward in the hierarchy so that those that did not correlate at level 1 may be correlated at level 2 as part of a larger collection (from a larger region of the network). This process continues upward until the anomalous reports aggregate at the root. Note that if the network has partitioned, a root will be selected in each partition and will become the aggregation point for alerts and anomalous path data.

3.4 Initial Distributed Detector Prototype

To investigate these ideas in more detail, we constructed a rudimentary wormhole creation tool and an initial distributed detector prototype for self-contained, reliable tunnel wormholes based on the approach discussed above. We installed these in a 19-node testbed that utilizes an enhanced version of the *Topology Scenario Manager* (TSM) from Telcordia Technologies' *Tactical Environment Assurance Laboratory* [LIT05] to emulate MANET topology, mobility, probabilistic packet loss (PPL). In tests using self-contained wormholes with reliable tunnels, 3% PPL, 60 randomly-generated topologies, and node mobility, the initial prototype successfully detected and localized the wormhole tunnel endpoints nodes without false alarms [KS05].

4. FAULT LOCALIZATION TECHNIQUE

Fault localization refers to the task of determining the location of a failure in a network; it is typically a difficult task because a failure commonly manifests itself in numerous symptoms at widely varying locations in the network. Many algorithms have been developed in the past for localizing a fault through correlation of the symptoms (also called alarms) received by a manager [SS04a, SS04b]. These algorithms vary in the network

models used, complexity of the computation, the assumptions made about the underlying network, etc. A recent promising new algorithm has been developed called Incremental Hypothesis Updating (IHU) which processes symptoms one at a time in an incremental fashion, thereby providing increased efficiency [SS04c]. This algorithm uses a probabilistic dependency model to represent the causal relationships between failures and symptoms along with probabilities that the symptoms are caused by the various failures. When symptoms are received, their associated failures in the dependency model are used to construct a set of hypotheses that can explain the causes of all the symptoms that have been observed. This hypothesis set is updated for each received symptom. The IHU algorithm has been shown to be fast, scalable, and accurate, with the potential of being deployable in real-time.

The dynamically changing dependencies of a battlefield MANET network cause difficulties for traditional fault localization techniques that use fixed dependency models because mobility may result in changed symptom-fault relationships causing an incorrect localization result. We have extended the IHU algorithm and its fault localization framework to an adaptive architecture that uses a dynamic dependency model, as shown in Figure 5 [NS05]. The dynamic dependency model used in this architecture incorporates temporal information into the fault-symptom relationships; the IHU algorithm is also modified to use this information in the construction of the fault hypotheses. The algorithm uses the dependency model to process the observed symptoms incrementally as they are received, and modifies the hypothesis on receiving the changed topology information. Incorporating temporal information into the localization process improves detection rates because the correct symptom-fault relationships are used in the analysis, and also helps distinguish the effects of mobility from real failures.




Figure 5: Adaptive fault localization system architecture

A promising new direction of research is active probing in which managers send probes to network nodes to diagnose network health [NS06]. These probes are test transactions whose outcome depends on the proper functioning of various network components. Fault localization using active probing involves two steps: problem detection in which probes are used to detect the presence of a failure, and problem determination in which additional probes are sent to determine the exact location of the failure.

5. INTEGRATED INTRUSION DETECTION AND FAULT LOCALIZATION

There are many similarities between the tasks of intrusion detection and fault localization. In both problem domains, we wish to determine the cause and location of an abnormality that causes the network to malfunction, and in both domains, this localization relies on the correlation of symptoms caused by the malfunction. The difference is that the malfunction is due to natural causes in fault localization whereas it is due to deliberate malicious interference in intrusion detection. The problem of fault localization may be somewhat easier because we can rely on probabilistic and statistical relationships between faults and symptoms and we are not dealing with an intelligent adversary who can modify his behavior or mount counter-measures to thwart detection. However, as illustrated in the wormhole example above, statistical knowledge of normal network conditions can help to detect an intrusion, construction of hypotheses similar to those in the IHU algorithm can help to narrow down possible locations of the attacker, and active probing techniques employed in narrow carefully chosen locations similar to those used in problem determination can complete the solution.

5.1 Integrated Approach

In the integrated approach, each network node is responsible for monitoring various performance measures associated with the traffic transmitted by the node. These include round-trip delay and loss rates to various destinations. The set of destinations that are monitored can vary dynamically depending on the current network topology (as obtained from the routing tables).

As described in Section 4, the network can be monitored by sending periodic probes to collect information about various network parameters. These probes could be simple probes like pings or could be more sophisticated and customized probes. The network parameters could be end-to-end delays, loss percentage, etc. A probe set

should be built to provide adequate coverage of the network but without excessive traffic overhead. Various approaches could be adopted to build a probe set. One possible approach could be to send pings from each node to its 3-hop neighbors.

Each node also monitors its own queuing delay for outgoing packets. When the performance statistics point to certain anomalous situations, a symptom report is generated and sent to the fault localization system. For instance, the onset of a wormhole tunnel may cause packet end-to-end delays and/or loss rates to increase on the routes that pass through the tunnel, and this can be detected and reported to the localization system.

Fault localization uses a dependency model that is constructed from topology information obtained from the local routing tables and which associates symptoms such as excessive delays and loss rates on an end-to-end path with possible failure conditions of the nodes that lie on that path. When a symptom is received from another node, some pre-processing of the symptom may be performed to eliminate false symptoms with the help of the higher degree of global knowledge available from cooperating with the intrusion detection system.

5.2 Example: Improving Wormhole Detection Accuracy By Integrating Fault Localization

For instance, the following two observations can be exploited to detect a wormhole:

1. Under healthy network conditions, the end-to-end delay could be explained by the sum of queuing delays on the hops of the end-to-end paths. During a wormhole attack, for some paths, the perceived path would be different from the actual path. On such paths, end-to-end delay will not be explained by the queuing delays on the hops of the perceived path.
2. Secondly, the queuing delay on a network node can be explained by the traffic going through that node. Under healthy network conditions, an increase in queuing delay could be explained by an increase in the network traffic observed on that node. However on the nodes that form a tunnel, the perceived traffic going through these nodes cannot explain an increase in the queuing delay on such nodes.

Observation 1 will hold for all paths that pretend to go directly from one tunnel end-point to another, but in reality go through the wormhole tunnel. Correlating such symptoms through a fault localization algorithm such as Incremental Hypothesis Updating (IHU) can lead us to identify the possible tunnel end-points. On the other hand, observation 2 will be seen on paths that do not get attracted by the wormhole but go through tunnel nodes.

For such paths, the perceived and actual hops would be same but the queuing delay on the hops that form the tunnel would not be explained by the end-to-end traffic going through these nodes. Correlating these symptoms through the IHU algorithm can lead us to identification of wormhole tunnel nodes.

The incremental IHU algorithm would then come up with a hypothesis set along with confidence levels for each hypothesis. This hypothesis set should allow a more precise determination of the likely problems that cause the observed symptoms and also point to the possible locations where the offending nodes are located.

The adaptive fault localization architecture described earlier would play an important role in the task of intrusion detection, since the dependency model will change as the network topology changes. The activation of a wormhole tunnel results in a perceived link to be created, which changes the topology as observed by various nodes in the network. The adaptive architecture will make it possible to correlate these topology changes with the received symptom reports that indicate performance anomalies.

5.3 Improving Fault Localization By Incorporating Hierarchical Organization

While already an integral part of our intrusion detection technique (see Section 3.1), we believe that fault localization can be handled more efficiently by deploying managers for each cluster and dividing the correlation task into intra-cluster and inter-cluster tasks. A *cluster-manager* collects per-node statistics and infers the anomalies between the end-to-end and per node statistics within its cluster. Such an approach requires a hierarchical organization of the management system in which network clusters are managed by separate managers. Managers of these lower level clusters report to their upper level cluster managers.

In this organization, the dependency model of the entire network is distributed among *domains*. These domains are defined by the clustering information. The dependency model built by each domain includes symptoms of those end-to-end paths that are entirely located in that domain. Communication with sub-domains may be represented in the dependency model by proxy nodes which contain information about the inter-connectivity of these sub-domains. Similarly the faults located outside the domain which may cause a symptom to arise within a domain may also be represented by proxy nodes. Thus the dependency model for a domain will have the same structure as in the centralized approach, but its scope will be smaller and interpretation of some nodes may be different.

A fault localization algorithm structured in this hierarchical manner relies on cooperation among domain managers. On observing an anomaly, the domain manager begins diagnosis of an end-to-end anomaly symptom only if all the nodes on that path are located within the domain. Otherwise, the symptom is delegated to the higher-level manager. While analyzing the failure of an end-to-end path completely located in its domain, the domain manager divides the diagnosis into various path segment diagnoses. It then delegates the task of diagnosis of a path segment to the manager of the sub-domain that entirely contains that path segment. Each upper level manager in this hierarchy then correlates the information collected from the diagnosis of its lower level managers using its dependency model.

6. CONCLUSIONS*

In this paper, we have proposed that intrusion detection and fault localization techniques in MANET environments should work cooperatively due to their many similarities. Not only would this integrated approach improve detection and localization accuracy, but it would also minimize system redundancy and reduce overhead by eliminating duplicate measurement probes. While our research is at an early stage, by using the detection of in-band wormholes as an example, we have outlined how an integrated approach can better distinguish malicious network attacks from “normal” network delays and outages. In future work, we plan to extend the fault localization hypotheses model to incorporate intrusion detection models and validate our ideas by simulation and experimentation.

REFERENCES

- [BK01] S Banerjee and S Khuller: "A Clustering Scheme for Hierarchical Control in Wireless Networks", Technical Report CS-TR-4103, Department of Computer Science, University of Maryland, College Park, 2001.
- [GA04] J. Garstka and D. Alberts, "Network Centric Operations Conceptual Framework", Version 2.0, June 2004, available from web site : <http://www.oft.osd.mil/initiatives/ncw/ncw.cfm>
- [HU03] Y-C. Hu, A. Perrig, and D.B. Johnson., "Packet Leashes: A Defense against Wormhole Attacks in Wireless Ad Hoc Networks", Proc. IEEE Infocomm 2003, San Francisco CA, March 2003.
- [KS05] P. Kruus, D. Sterne, et al, "In-Band Wormholes and Countermeasures in OLSR Networks", to appear in Proc. SecureComm 2006, Baltimore, MD, August 2006.
- [LIT05] M. Little, "TEALab: A Testbed for Ad Hoc Networking Security Research", IEEE Military Communications Conference, Atlantic City, NJ, October 2005.
- [MMM06] K. Manousakis, A.J. McAuley, et al, "Creating and Maintaining a Good Intrusion Detection Hierarchy in Dynamic Ad Hoc Networks", submitted to MILCOM 2006, Washington DC, October 2006.
- [NS05] M. Natsu and A.S. Sethi, "Adaptive Fault Localization for Mobile, Ad-Hoc Battlefield Networks." Proc. Milcom-2005, IEEE Military Communications Conference, Atlantic City, NJ, October 2005.
- [NS06] M. Natsu and A.S. Sethi, "Active Probing Approach for Fault Localization in Computer Networks." Proc. E2EMON-2006, IFIP/IEEE End-to-End Monitoring Workshop, Vancouver, B.C., Canada, April 2006.
- [OLSR03] "Optimized Link State Routing (OLSR)," IETF RFC 3626, T. Clausen, P. Jacquet, Ed., October 2003.
- [SBC05] D. Sterne, P. Balasubramanyam, D. Carman, B. Wilson, R. Talpade, C. Ko, R. Balupari, C-Y Tseng, T. Bowen, K. Levitt, J. Rowe, "A General Cooperative Intrusion Detection Architecture for MANETs", Proc. Third IEEE International Information Assurance Workshop, College Park MD, March 2005.
- [SS04a] M. Steinder and A.S. Sethi, "Non-deterministic fault localization in communication systems using belief networks." *IEEE/ACM Transactions on Networking*, Vol. 12 No. 5, pp. 809-822, October 2004.
- [SS04b] M. Steinder and A.S. Sethi, "A survey of fault localization techniques in computer networks". *Science of Computer Programming, Special Edition on Topics in System Administration*, Vol. 53 No. 2, pp. 165-194, November 2004.
- [SS04c] M. Steinder and A. S. Sethi, "Probabilistic fault diagnosis in communication systems through incremental hypothesis updating". *Computer Networks*, Vol. 45 No. 4, pp. 537-562, July 2004.

* The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.