

Coordinated Packet Traceback in Wireless Mesh Networks

Lichun Bao and Denh Sy

1 Introduction

Wireless Mesh Networks (WMNs) have emerged as an important technology to deliver pervasive network access and ubiquitous mobile applications in various scenarios, such as disaster relief operations, ad hoc network support in convention centers, remote or hard-to-reach terrains with their promise of rapid deployments and reconfigurable architecture. In most cases, WMNs were intended to provide Internet access capabilities by extending existing network coverage areas, and reducing the cost of backhaul connections to distant base-stations. Standard organizations are actively calling for specifications for mesh networking, for instances, IEEE 802.11, IEEE 802.15, IEEE 802.16 and IEEE 802.20.

Security is one of the most important issues in WMNs to ensure the trusted communications. Besides robust networking protocol design requirements to combat security exploitations, application-level intrusion detection and traceback mechanisms are necessary in many situations. A particularly harmful attack against WMNs is known as the Sybil attack [10], where a node illegitimately claims multiple identities. Without bulletproof hardwares, packet traceback seems to be the only mechanism left to track down Sybil attackers.

Unfortunately, a lot of the existing research in the packet traceback techniques and architectures were designed for Internet packet traceback, and are not viable in WMNs because of the characteristics of WMNs. First of all, human intervention or in-network intrusion detection are not possible to prevent attacks due to the wide deployment nature in WMNs for distributed processing, and the limited power and silicon supply. Secondly, the wireless transmission medium is openly accessible to anyone in the network adjacency, thus making the network vulnerable to malicious packet injections. Third, WMNs are capable of altering packet information in many peer-to-peer applications, which makes address-based packet filtering and location inference impossible.

Fortunately, what makes WMNs vulnerable is also the strength of WMNs for security functions. We propose CAPTRA, a CoordinAted Packet TRAcEback protocol, for WMNs that takes advantage of the open transmission and in-network processing power of the WMNs, and uses multi-dimensional Bloom filters for packet traceback purposes. In CAPTRA, each mesh node allocates a small memory block for the Bloom filter. When a packet of interested categories traverses the network, each forwarding mesh node records the packet in its Bloom filter, so do the overhearing nodes. Different from previous approaches that build Bloom filter hash functions based on the packet information alone, the forwarding node ID is taken as well an input to the

hash function in the Bloom filter in CAPTRA. During traceback operations, the predecessor on the packet source route is determined by a collaborative mechanism by summarizing all possible packet witnesses that are collected from the overhearing nodes. If the witness quorum reaches certain threshold, traceback operations are recursively continued toward the true origin of the packet hop-by-hop. Three messages, REQUEST, VERDICT and CONFESS, are used for the traceback coordination, witness collection, and source route construction processes, respectively. To our knowledge, packet traceback in WMNs is among the few with little research so far.

2 Related Work

Network attacks can be either persistent or sporadic [28]. In persistent attacks, the offenders must frequently launch attack packets to bombard the victims, which can be relatively easy to trace by mechanisms such as probabilistic packet marking, traffic logging, data mining etc. Whereas in sporadic attacks, a single packet can render havoc at the potential victim, such as the WinNuke, Ping of Death and Teardrop attacks.

Intrusion Detection Systems (IDSs) use *attack signature* or *pattern* to help distinguish malicious packets from normal traffic. At the very least, an attack signature is defined by the IP address or address range of the entity that is being attacked. A variety of methods for IDSs were discussed in [26] that can help detect if received packets have spoofed source addresses.

After malicious attacks are detected, the subsequent traceback to the attack origins requires the network participate in the pre-attack tracking and post-attack tracing operations, in which *packet tracking* refers to the recording of packets or packet flows when the packets are forwarded from their sources to their destinations, and *packet tracing* refers to the operations to find out the *source route* of a packet, which is the sequence of nodes that have forwarded the packet or the flow. This term is slightly different from what is commonly known by the source routing in IP forwarding.

Source route identification problem is commonly referred to as *network traceback*, or *IP packet traceback* in the Internet arena. Two network tracing problems are currently being studied: “single-domain IP traceback” and “traceback across stepping-stones” (or a “connection chain”). Traceback across stepping-stones is to identify the origin of an anonymous attacker through a chain of connections before the attacker interacts with a victim host. The Decentralized Source Identification System (DECIDUOUS) uses IPsec security associations (SAs) and authentication headers to dynamically deploy secure authentication tunnels, and traces back to the attacks origins [5, 6]. The premise of DECIDUOUS is that if an attack packet has been correctly authenticated by a certain router, the attack packet must have been transmitted through that router. It utilizes IPsec security associations to dynamically deploy secure authentication tunnels in order to further trace down the possible attackers’ locations. Other approaches such as correlating the inter-packet delay were also proposed [30].

Many approaches have been proposed to trace single-domain IP traceback. The IP marking approaches enable routers to probabilistically mark packets with partial path information, and try to reconstruct the complete path from the packets with the markings [12, 17, 22, 24, 29].

A more explicit approach using IP marking is to add the IP address of the router to the IP header by turning on the *Record Route* option of the IP header [9]. An algebraic approach is proposed to transform the IP traceback problem into a polynomial reconstruction problem, and uses techniques from algebraic coding theory to recover the true origin of spoofed IP packets by having routers to embed information randomly into packets [8]. This is similar to the technique used in [22], which uses algebraic techniques to encode the path information as points on polynomials, and then reconstruct these polynomials at the victim. Li *et al.* proposed a hybrid approach to selective mark packets for tracking in the intermediate routers using Bloom filters [14].

ICMP traceback (iTrace) proposes to introduce a new ICMP message (or an iTrace Traceback message) so that routers can, with a low probability, generate iTrace messages to help the victim or its upstream ISP to identify the source of spoofed IP packets [2]. With enough ICMP Traceback messages from enough routers along the path, the traffic source and path can be determined. An intention-driven iTrace is also introduced to reduce unnecessary iTrace messages to improve the performance of iTrace systems [15].

An IP overlay network-based traceback system, named CenterTrack, selectively reroutes interesting IP packets directly from edge routers to special tracing routers, and the hop-by-hop input-debugging is then used [25]. *Input debugging* refers to the diagnostic features required to recursively determine from which adjacency a packet arrived that matches an attack signature on an individual router, until the edge of the network is reached and the edge ingress adjacency is identified.

Snoeren *et al.* proposed an architecture, Source Path Isolation Engine (SPIE), that integrates the IDS and single-packet traceback engines [23] to reconstruct the identify the attack graph. According to SPIE, once the IDS detects an abnormal attack event, the attack packet is fed into the traceback manager to generate a traceback request, which is sent to multiple network agents for constructing the regional attack graph based on the attack packet. Afterward, the regional attack graphs are assembled into a complete attack graph at the traceback manager, and fed back the IDS. SPIE is based on Bloom filter [3] for packet logging purposes [21].

3 Multi-Dimensional Hash Table

3.1 Bloom Filter and Its Variants

A Bloom filter is a space-efficient probabilistic data structure that is used to test whether or not an element is a member of a set [3]. Bloom filters are used in myriad of applications. Wherever a list or set is used, and space is a consideration, a Bloom filter is commonly considered [4].

Usually, elements are only added to the set, but not removed. Given a finite set, *false positive* judgment of the membership are possible, but *false negative* judgment are not in the traditional Bloom filters without refreshing. When the more elements are added to the set, the probability of false positive becomes greater.

Several variants of Bloom filters have been proposed. Attenuated Bloom filters [18] use arrays of Bloom filters to store shortest path distance information. Spectral Bloom filters [20] extend the data structure to support estimates of frequencies. In Counting Bloom Filters [11] each entry in the filter need not be a single bit but rather a small counter. Insertions and deletions to the filter increment or decrement the counters respectively. When the filter is intended to be passed as a message, compressed Bloom filters [16] may be used, where parameters can be adjusted to the desired tradeoff between size and false-positive rate.

Space-code Bloom filter (SCBF) provides frequency estimation of an element by probabilistically filling up multiple normal Bloom filters, from which to statistically infer the frequency of the element [13]. In SCBF, a randomly chosen Bloom filter of the SCBF module takes the flow ID (source and destination IP addresses and port numbers) as the inputs, and records the membership in the table. When the Bloom filters in the SCBF module saturate, the contents of the Bloom filters are paged out into a log file. Later, according to the number of positive answers to a flow query into the log files offline, the flow volume can be probabilistically inferred.

3.2 Multi-Dimensional Hash Table

It is easy to see that a Bloom filter can be reduced to a hash table based on a single hash function, where $k = 1$. Alternatively, a Bloom filter is augmentation of a single hash table by the multiplicity of hashing functions. We generalize the construction of Bloom filters by introducing the concept of *dimensions* in hash algorithms, in which a dimension can expand by the number of either hash functions, hash tables, or both. We count the instances of expansions, and denote the number of dimensions accordingly.

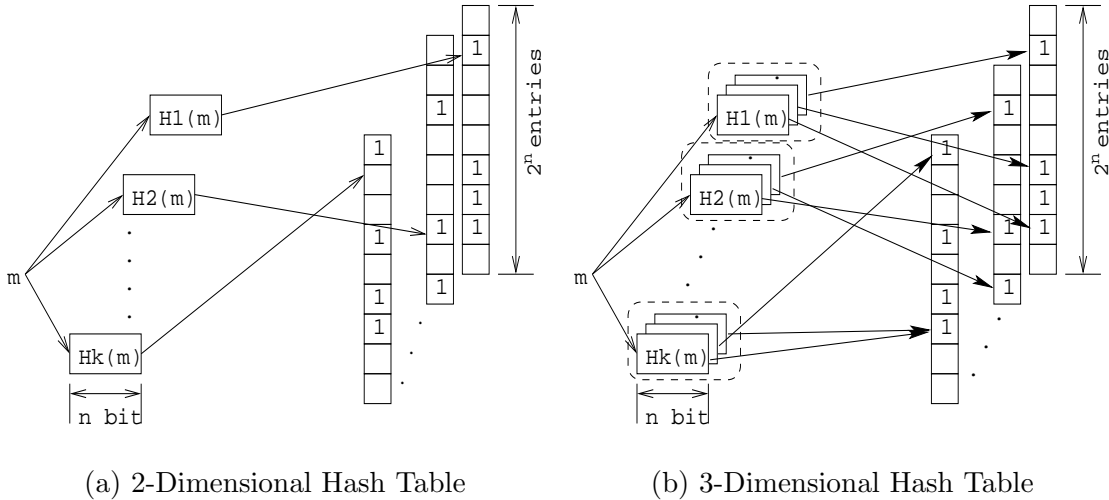


Figure 1: Dimensions of Hash Functions in Filters

For instances, a hash table based on a single hashing operation is one-dimensional. The Bloom filter is a 2-dimensional hash table because a Bloom filter is implemented using k hashing functions in an array. In addition, the Bloom filter table can be split into k separate hash tables,

respectively, as shown in Figure 1(a), which is another 2-dimensional hash table. In addition, the individual hash functions in Figure 1(a) can again be augmented by k different hash functions, thus making it a 3-dimensional hash table, as shown in Figure 1(b).

In particular, we refer to the 3-dimensional hash table in Figure 1(b) as *Space-Time Bloom Filter* (STBF) in CAPTRA for WMN packet traceback because the 2-dimensional Bloom filter data structure is replicated and scattered among multiple mesh nodes, and updated asynchronously for packet tracking in different parts of the WMN. It is straightforward that the aforementioned space-code Bloom filter (SCBF) is another 3-dimensional Bloom filter in Figure 1(b) for a different purpose.

By defining the dimensions in hash tables, we have generalize the concept of Bloom filter to a simpler but more comprehensive form, and captures all the hash algorithms in this chapter. Naturally, the dimension of hashing operation is unlimited. Depending on how each Bloom filter in the dimensions is utilized, we can apply the appropriate dimensions of hashing operations for membership query, such as packet tracing, routing and accounting purposes [13].

4 Coordinated Packet Traceback (CAPTRA)

We assume that a wireless mesh network (WMN) consists of simple networking computers that are interconnected via a shared wireless channel using the same channel access control protocol, such as the IEEE 802.11 [1]. In many WMN deployments, the use of access points (APs) as an aggregation point and the entrance to the Internet is a common practice. APs possess sufficient computing power to implement the intrusion detection mechanisms and initiate traceback operations. In addition, we assume that the mesh nodes' memories are limited so that the Bloom filters cannot be arbitrarily large, nor can there be permanent storage for logging purposes.

The choice of the Bloom filter algorithm is a local choice dependent on the memory resources and computational tasks of the mesh node, and its parameters on each mesh node are configurable, such as the table sizes. The reporting of a hit in the Bloom filter at a wireless node is also dependent on the capability of the node. Nodes in relatively static and dense wireless environments can report the hits less frequently, because the voting quorum can easily be satisfied, than nodes in highly mobile and sparse networks for the same level of traceback accuracy.

For simplicity, we assume that WMN nodes are homogeneous in terms of memory allocation, hashing algorithm and hit reporting frequency.

The essential difference of packet traceback in wireless networks from wired networks operations is that the transmission medium is open and can be overheard by any node in the nearby vicinity of the packet transmitter. Therefore, we can take advantage of such fact in the Space-Time Bloom Filter constructions. That is, the Bloom filters of a single Space-Time Bloom Filter are virtually dispersed among adjacent wireless mesh nodes, and the Bloom filter lookup is now a majority-vote by all the Bloom filters of the potential nodes.

The distribution of the virtual STBF is especially useful in WMNs because a) the aggregated Space-Time Bloom Filter provides larger memory capacity and parallel computing power than a

single node; b) the traffic intensity is uneven in different locations of the WMNs, thus distributing the Bloom filters off-loads the memory requirements of the nodes with high traffic in its neighborhood; c) nodal mobility often invalidates the source routes constructed from the packet traceback. By distributing the Space-Time Bloom Filters to multiple nodes in a vicinity, the traceback operations can tolerate certain node losses along the source route because the predecessor of the missing node can still be detected if there is enough number of witnesses to the packet being sent from the predecessor.

4.1 Hash Functions

Because the main purpose of packet traceback is to capture the path traversed by a specific packet through the network, we extract the certain pieces of information to identify the packet. For instance, the identification of an IP packet may include IP version, header length, source and destination addresses, fragment information and portions of the payload [23], which is then concatenated and packed into a bit string as the input to hash functions.

In order to use the same hash algorithm for different hash functions in the Space-Time Bloom filter (STBF), we feed additional information to the hash functions for packet tracing purpose:

1. An identifier for each hash function, which includes two pieces of information: (a) the host ID, and (b) a unique ID for the hash function.
2. The predecessor from which the packet is received.

By including the hash function ID, we can use the same hashing algorithm, such as MD5 [19] to compose different hash functions in the Space-Time Bloom Filter. By including the predecessor ID, a Bloom filter hit for a packet also reveals the previous hop of the packet. In the actual computations, all neighbor IDs have to be checked in the Bloom filter in order to find out the predecessor.

4.2 Query in the Space-Time Bloom Filter

As an STBF continuously monitors network traffic, the hash tables can gradually saturate, and the probability of false positive increases. To guarantee the best performance tradeoff between the memory utilization and false positive rate, a “50% Golden Rule” is derived and applied so that the Bloom filter can hold the maximum number of elements [3, 13] with lowest false positive rate.

Previously applications of Bloom filters for accounting or tracing purposes reset the whole Bloom filter memory when the filter saturation crosses the threshold. Therefore, the false positive rate is a variable, while the false negative rate is either 0 before flushing or 1 after flushing for certain elements.

In the Space-Time Bloom Filter (STBF), individual Bloom filter elements are dispersed among multiple nodes. When the saturation ratio of an individual hash table crosses the 50% threshold,

we only wipe out that single hash table in the STBF. Therefore, we define a majority-vote principle for STBF that allows STBF to yield a “hit” response when the number of “hit” responses from different Bloom filter elements satisfies a quorum. This is drastically different from the membership query operation in the traditional Bloom filter applications, where a hit requires unanimous hits on all the hash functions.

Due to the reset function on the STBF, the majority-vote quorum is a trade-off between the false negative and false positive rates in STBFs. If the quorum is low, the false positive rate increases, and if high, the false negative rate increases.

Intuitively, the majority-vote mechanism in STBF maintains gradually fading memory of the packet events in the past, and give more acute memory to recent traffic flows. We see increasing probability of false negatives because the individual hash tables are asynchronously being reset. In contrast, previous Bloom filter resets abruptly change the false positive and false negative rates of the Bloom filter to 0 or 1 in previous applications.

In order to avoid possible synchronized hash table resets, we add in certain randomness to the reset timing when the saturation rate reaches the threshold.

4.3 Traceback Messages

In support of packet traceback operations, three traceback messages are used, which jointly provide the majority-vote mechanism to corroborate the conviction of a node as one on the packet source route.

Message TRAC_{REQ} is to initiate a traceback query to discover the predecessor of a packet, used for ping-ing a network host. It is started by the access point in WMNs, and recursively carried out by the nodes on the packet source route. In the payload field of the traceback request message, the packet identification is carried.

Upon each message TRAC_{REQ} arrival, the packet traceback protocol extracts the packet identification, and runs the hash functions on all the current one-hop neighbors to determine the potential predecessor of the packet. If such predecessor exists, the traceback operation can continue by sending the traceback request to the predecessor. Message TRAC_{VERD} is for a node to indicate that it has witnessed the packet, and to issue a verdict to the sending node of the packet.

Message TRAC_{CONF} is sent to the WMN access point by a node c to tell that node c has forwarded the packet before if it has collected enough verdicts from its neighbors. In addition, node c propagate the traceback request farther to its candidate neighbors for farther traceback.

4.4 Packet Tracking

In CAPTRA, each node maintains a Bloom filter using k hashing functions. As mentioned before, upon overhearing a packet from a sender, a node extracts a bit stream from the packet, and hashes it into the node’s Bloom filter. The bit stream concatenates the packet identification information, the sender’s ID and the node’s ID. For instance, if node i hears a packet pkt with

identification information pkt.id , the packet is tracked by a set of k bits in node i 's Bloom filter, which are computed by $\text{Hash}(\text{pkt.id-pkt.tx-}i\text{-}j)$, in which Hash is a hash function, pkt.id is the packet identification information, pkt.tx is the transmitter ID of the packet, the symbol “-” represents the concatenation of the operands, k is the number of hash functions in the Bloom filter, and $j = 1 \dots k$.

A packet is tracked when a node forwards the packet, or a node overhears and successfully receives the packet. Because the overhearing node's ID is involved in the hashing algorithm, each node overhearing the packet maps the packet to a different Bloom filter entry. Therefore, the Bloom filter fill-up rates are different at different nodes.

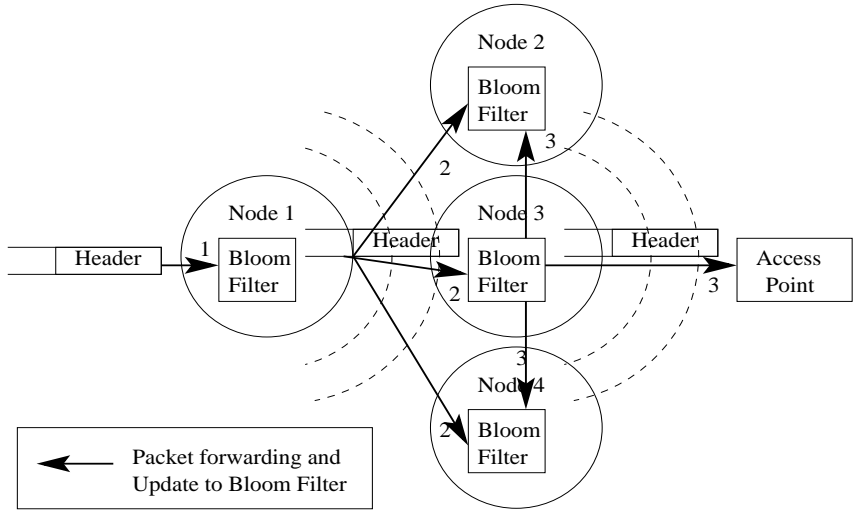


Figure 2: Space-Time Bloom Filter for Packet Tracking

Figure 2 illustrates a partial WMN with three mesh nodes and an access point for data collection purposes. The blue arrows indicate the packet transmissions and receptions. The senders and the receivers of the packet record the packet in their Bloom filters.

According to the “50% Golden Rule”, the Bloom filters are refreshed once the saturation ratio reaches 50% of the Bloom filter capacity, and every bit of the Bloom filter is reset to 0.

4.5 Packet Tracing

In order to describe the packet traceback operations in CAPTRA, we again look the partial network in Figure 2 as an example for illustrations.

Suppose that, in Figure 2, the access point initiates a traceback query on packet pkt due to a security breach detected by an IDS. The access point sends out the traceback request with the packet identification pkt.id in the newly defined message TRAC_{REQ} with pkt.id as payload. Because the access point can look up its own Bloom filter for a hit, and find out the sender of the packet identified by pkt.id , the traceback request is sent directly to the packet predecessor — node 3, and is overheard by the nodes around the access point.

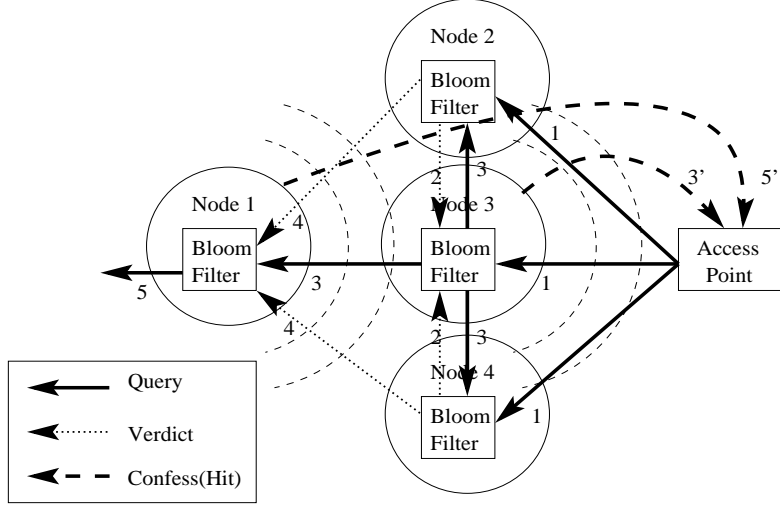


Figure 3: Packet Tracing Process using Space-Time Bloom Filter

Due to the broadcast nature of the wireless medium, nodes 2 and 4 overhear the traceback request in Figure 3. Because nodes 2 and 4 overheard about packet pkt when it was forwarded 2, they must have recorded pkt.id plus the sender ID in their Bloom filters. If there is a positive hit, node 2 or 4 sends the $\text{TRAC}_{\text{VERD}}$ packet to node 3 to reinforce the tracing request.

When enough verdicts are collected after a period of time, including the Bloom filter hits at nodes 2, 3 and 4, and the access point, node 3 sends back a $\text{TRAC}_{\text{CONF}}$ packet to the access point for attack source route construction. The threshold of the number of verdicts for further traceback is a tunable parameter, and is at least 2 because the traceback request sender and receiver also generate verdicts. Note that each traceback request is a verdict to the receiver of the traceback request.

In addition, node 3 generates further traceback request $\text{TRAC}_{\text{REQ}}(\text{pkt.id})$ based on its current neighbors, and in Figure 3, a new request is generated to node 1, after which the verdict generation process follows the same as above. If in rare cases, multiple one-hop neighbors of node 3 appear to be the predecessor of the packet, each of them receives the traceback request. If a node is wrongly convicted of the packet transfer, it is expected the wrong traceback request gradually die out because of lack of evidence further away from the convicted node.

4.6 Loop Avoidance in Packet Traceback

Without addition mechanisms, the above traceback protocol can cause loops while forwarding the traceback request, one instance in Figure 3 being that node 1 sends the traceback request again to node 3 because node 1 has heard about node 3's transmission of the packet, which causes unnecessary traceback activities.

To avoid the potential looping problem, the traceback request is appended with distance information, indicating the distance from the current sender of traceback request to the origin of the traceback request.

If a node receiving the traceback request has a shorter distance to the request origin than the sender of the request, the traceback request is discarded. Note that the criteria based on the reverse path to the origin may be incorrect due to network mobility.

4.7 Termination of Tracing Query

When a node has no other one-hop suspect but itself who has transfer the packet, the tracing ends.

Some optimizations are possible in CAPTRA. In case the suspect which receives the traceback request is an uncooperative node, and refuses to respond to the request for packet tracing, the request generator may be able to conclude that the suspect has actually transfer the packet according to the verdicts from their shared one-hop neighbors, *e.g.*, nodes 1 and 3 in Figure 3.

5 Simulation Evaluations

We used J-Sim [27] to simulate the packet traceback protocol in WMNs. J-Sim is a network simulator constructed entirely in JAVA.

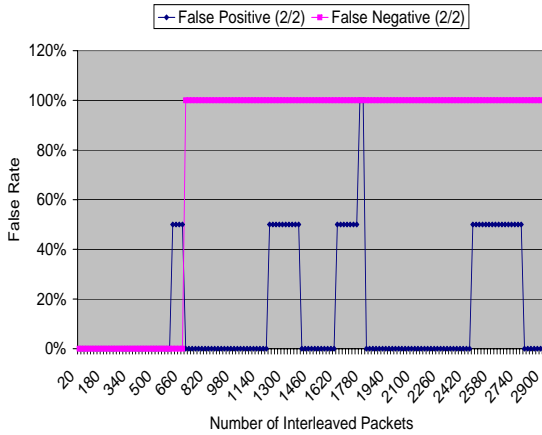
In the simulations, twenty wireless nodes, numbered from 0 to 19, are deployed in a linear fashion, spaced 8 meters apart. Each node has a 25-meter transmission range, so that each node has up to 7 neighbors. The propagation and path-loss model use the free-space model. AODV is used as the underlying routing protocol. Each node contains a Bloom filter of size 4096 bits, which is refreshed using the “50% Golden Rule”. The number of hash functions for each packet, k , is a system-wide variable, and varies in two simulation scenarios from 2 to 3 for comparison purposes.

The traffic source is attached to node 0, and destined to node 19. A CBR traffic generation model with 512 byte per packet at a rate of 10 packets per second was used. We change the location node 0 in the network so as to change the hop distance to node 19 in order evaluate the performance packet tracing with regard to the hop distance.


In each simulation round, the packet tracking capacity of the network is tested by collecting the false positive and false negative rates when different amount of traffic is forwarded between the times when the packet is generated at the source and when the packet is being traced back from node 19. As discussed earlier, the CAPTRA voting system is used to convict a node, and have a node confess. In the simulations, a node confesses as a packet forwarder when two nodes convicts it.

5.1 Simulation Results and Analysis

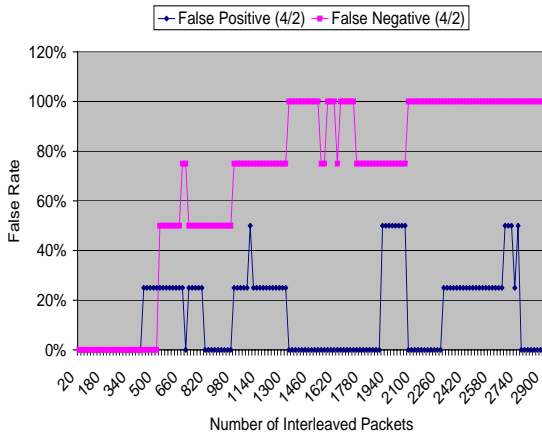
Figure 4 presents the false positive and false negative rates under the various simulation scenarios. Two types of Bloom filters are implemented and compared side-by-side, in which one type of Bloom filter uses 2 hash functions in all the simulation scenarios, and the other uses 3 hash functions in each. The false positive rates are simply the number of false positives divided by the



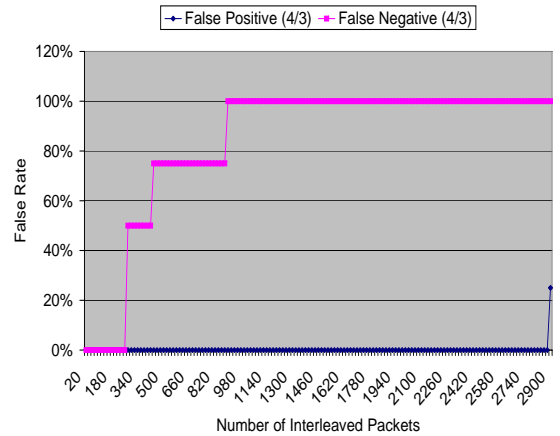
(a) 2-hop Traceback with 2 Hash Keys



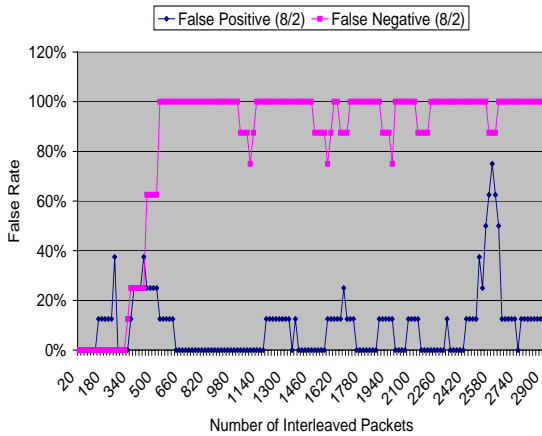
(b) 2-hop Traceback with 3 Hash Keys




(c) 4-hop Traceback with 2 Hash Keys



(d) 4-hop Traceback with 3 Hash Keys



(e) 8-hop Traceback with 2 Hash Keys



(f) 8-hop Traceback with 3 Hash Keys

Figure 4: The False Positive and Negative Rates in Multihop Packet Traceback

hop count in the WMN. Similarly, false negative rates are the number of false negatives divided by the hop count.

In Figure 4 (a) (c) and (e), the traceback performance is measured for 2-, 4- and 8-hop source route of a particular packet. As we can see, when the number of hops increases, the network has shorter and shorter memory of the traced packet, indicated by the less number of interleaved packets from the time when the traced packet is generated to the time when the packet is traced back. In addition, the duration of the network memory about the packet is depended on the network density, and the size of the Bloom filters, which are 4096 bits in our simulations.

On the other hand, Figure 4 (b) (d) and (f) show the false positive and false negative rates of similar simulations when the number of hash functions is three. Comparing with the other three corresponding sub-diagrams, these three settings presents more stable traceback performance, but shorter memory of the traced packet due to the faster fill-up rate of the Bloom filters. Therefore, the choice of Bloom filters depends on the application trade-offs between accuracy of traceback and the duration of the traceback validity.

There are unstable false rates on the curves of the false positive and false negative rates. This is due to the Bloom filter collisions. Due to space limitations, we skipped various other settings that variate and mix different Bloom filter setup strategies.

6 Summary

We have introduced a new perspective on expanding Bloom filter dimensions, which led to the Space-Time Bloom Filter (STBF) for coordinated packet traceback (CAPTRA) in wireless mesh networks (WMNs). CAPTRA takes advantage of the broadcast nature of the wireless transmissions for coordinated traceback, and is especially suitable for WMNs due to the small memory and computational requirements in maintaining STBF. The space-time Bloom filter construction enables distributed maintenance and asynchronous and gradual refreshing of the Bloom filters so that the WMN keeps robust and gradually fading memory of the packet traversal event. These unique organization and utilization of Bloom filters in CAPTRA allows our future application of packet traceback in adaptive routing in WMNs.

References

- [1] IEEE Std 802.11. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Technical report, IEEE, Jul. 1997.
- [2] S.M. Bellovin, Marcus Leech, and Tom Taylor. ICMP Traceback Messages. Technical report, Internet Draft, IETF, Mar. 2001.
- [3] B.H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of ACM*, 13(7):422–426, Jul. 1970.
- [4] A. Broder and M. Mitzenmacher. Network applications of Bloom filters: a survey. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, 2002.
- [5] H.Y. Chang, P. Chen, A. Hayatnagarkar, R. Narayan, P. Sheth, N. Vo, C. L. Wu, S.F. Wu, L. Zhang, X. Zhang, F. Gong, F. Jou, C. Sargor, and X. Wu. Design and Implementation of A Real-Time Decentralized Source Identification System for Untrusted IP Packets. In *Proceedings of the DARPA Information Survivability Conference and Exposition*, Jan. 2000.

- [6] H.Y. Chang, R. Narayan, C. Sargor, F. Jou, S.F. Wu, B.M. Vetter, F. Gong, X. Wang, M. Brown, and J.J. Yuill. DECIDUOUS: Decentralized Source Identification for Network-Based Intrusions. In *Proceeding of 6th IFIP/IEEE International Symposium on Integrated Network Management*, pages 702–714, 1999.
- [7] Inc. CrossBow Technology. <http://www.xbow.com>, 2005.
- [8] D. Dean, M. Franklin, and A. Stubblefield. An Algebraic Approach to IP Traceback. In *Proceedings of Network and Distributed System Security Symposium*, Feb. 2001.
- [9] T.W. Doepfner, P. N. Klein, and A. Koyfman. Using Router Stamping to Identify the Source of IP Packets. In *7th ACM Conference on Computer and Communications Security*, pages 184–189, Athens, Greece, Nov. 2000.
- [10] J. R. Douceur. The Sybil attack. In *First International Workshop on Peer-to-Peer Systems (IPTPS)*, Mar. 2002.
- [11] L. Fan, P. Cao, J. Almeida, and A. Broder. Summary cache: A scalable wide-area Web cache sharing protocol. In *Proceeding of SIGCOMM*, 1998.
- [12] M.T. Goodrich. Efficient Packet Marking for Large-Scale IP Traceback. In *9th ACM Conf. on Computer and Communications Security (CCS)*, pages 117–126, 2002.
- [13] A. Kumar, J. Xu, E.L. Li, and J. Wang. Space-code bloom filter for efficient traffic flow measurement. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 167–172, Miami Beach, FL, 2003.
- [14] J. Li, M. Sung, J. Xu, and L. Li. Large-Scale IP Traceback in High-Speed Internet: Practical Techniques and Theoretical Foundation. In *IEEE Symposium on Security and Privacy*, Berkeley, CA, May 2004.
- [15] A. Mankin, D. Massey, C. Wu, S. F. Wu, and L. Zhang. On Design and Evaluation of Intention-Driven ICMP Traceback. In *Proceedings of IEEE International Conference on Computer Communications and Networks (IC3N)*, 2001.
- [16] M. Mitzenmacher. Compressed Bloom Filters. *IEEE/ACM Transactions on Networks*, 10(3):613–620, Oct. 2002.
- [17] K. Park and H. Lee. On the Effectiveness of Probabilistic Packet Marking for IP Traceback. In *Proceedings of SIGCOMM*, pages 15–26, 2001.
- [18] S.C. Rhea and J. Kubiawicz. Probabilistic Location and Routing. In *INFOCOM*, 2002.
- [19] R. Rivest. RFC 1321 - The MD5 Message-Digest Algorithm. Technical report, MIT Laboratory for Computer Science and RSA Data Security, Inc., Network Working Group, Apr. 1992.
- [20] Y. Matias S. Cohen. Spectral Bloom Filters. In *SIGMOD Conference on Management of Data*, pages 241–252, 2003.
- [21] L.A. Sanchez, W.C. Milliken, A.C. Snoeren, F. Tchakountio, C.E. Jones, S.T. Kent, C. Partridge, and W.T. Strayer. Hardware Support for a Hash-Based IP Traceback. In *Proceedings of DARPA Information Survivability Conference and Exposition*, Jun. 2001.
- [22] S. Savage, D. Wetherall, A. Karlin, , and T. Anderson. Practical Network Support for IP Traceback. In *Proceedings of ACM SIGCOMM Conference*, Aug. 2000.
- [23] A.C. Snoeren, C. Partridge, L.A. Sanchez, C.E. Jones, F. Tchakountio, S.T. Kent, and W.T. Strayer. Hash-based IP Traceback. In *Proceedings of ACM Conference on Applications, Technologies, Architectures and Protocols for Computer Communication (SIGCOMM)*, pages 3–14, 2001.
- [24] D.X. Song and A. Perrig. Advanced and Authenticated Marking Scheme for IP Traceback. In *Proceedings of IEEE INFOCOM Conference*, 2001.
- [25] R. Stone. CenterTrack: An IP Overlay Network for Tracking DoS Floods. In *Proceedings of 9th Usenix Security Symposium*, Aug. 2000.
- [26] S. Templeton and K. Levitt. Detecting spoofed packets. In *Proceedings of The Third DARPA Information Survivability Conference and Exposition (DISCEX)*, 2003.
- [27] H. Tyan. J-Sim. <http://www.j-sim.org/>.
- [28] B. Vetter, F. Wang, and S.F. Wu. An Experimental Study of Insider Attacks for the OSPF Routing Protocol. In *IEEE International Conference on Network Protocols (ICNP)*, pages 293–300, Oct. 1997.

- [29] M. Waldvogel. GOSSIB vs. IP Traceback Rumors. In *Proceedings of 18th Annual Computer Security Applications Conference (ACSAC)*, Dec. 2002.
- [30] X. Wang and D.S. Reeves. Robust correlation of encrypted attack traffic through stepping stones by manipulation of interpacket delays. In *ACM Conference on Computer and Communications Security*, pages 20–29, 2003.