Security in Different Management Domains of MANET

[1]Savita Bamal, [2]Archana Rohilla

[1, 2]Department of Computer Science and Engineering Gurgaon Institute of Technology and Management Gurgaon, Haryana, India
Email: [1]savita_nc@yahoo.com, [2] rohilla.archana@yahoo.com

Abstract : Abstract—Creating a division between inter-domain and intra domain routing is expected to assist in meeting the challenges of future MANET deployments in terms of heterogeneity and administrative separation. Previous research in MANET routing and security has focussed on the intra-domain context. This paper discusses the problems encountered in inter-domain routing security, and the security requirements that arise from interactions between multiple management domains. It identifies the ways in which inter-domain routing for MANETs requires a different trust model to that of intra-domain MANET routing. Additionally, it discusses the way in which its requirements differ from inter-domain routing security in the Internet due to the more dynamic topological environment found in MANETs. The paper identifies how domain ownership and relationships between domain owners influence the security problem.

Index Terms—Mobile Ad hoc Network, inter-domain routing, intra-domain routing.

I. INTRODUCTION

Mobile ad hoc networks (MANETs) have been developed to support communications in dynamic network topologies. However, most approaches to date have considered MANET routing from an intra-domain perspective with a flat network under a single administrative entity. In order to allow interoperation among heterogeneous network domains operated by different organisations, inter-domain routing approaches, such as IDRM (Inter-Domain Routing Protocol for MANETs) [4] are now being considered. This paper discusses the trust and key management models of previous work in the intra-domain MANET routing and inter-domain Internet routing spaces. It then identifies how the inter-domain MANET routing problem differs from these. Potential approaches to securing inter-MANET routing are then discussed.

II. CURRENT SECURITY MODELS

As a step towards designing security mechanisms for MANET inter-domain routing, for a protocol such as IDRM, we develop a model of actors in the network, their relationships, and the information that needs to be protected. Previous research in intra-domain MANET routing and inter-domain routing on the Internet provides some background. The intra-domain MANET routing problem requires support for dynamic network topologies, whilst the interdomain Internet routing problem provides a model of the entities and information components relevant to the provision of security across disparate management domains.

A. Intra-domain Manet Routing Security

It is a particular feature of many MANET routing scenarios that any device can act as a router, and the network can be rearranged in an arbitrary fashion. This means that assumptions about topology cannot be made when configuring security relationships between peer routers. From a security perspective, the primary questions are whether a router is authorised to participate in the network, and whether it owns a particular address. It is usually assumed that all the entities in the network are under a single administrative domain. The division between proactive and reactive protocols is reflected in the security mechanisms. SAODV [6][5] provides an example of reactive routing protocol security. It uses hash chains to avoid manipulation of hop counts in route discovery messages, and digital signatures are used for the immutable parts of these messages, to provide end-to-end confirmation that the request reached the owner of the address. SLSP [16] is an example of a security mechanism for a proactive routing protocol. It uses signatures on link state update messages to avoid manipulation of the topology information. The SAODV solution is focussed on verifying the validity of the path, whereas the SLSP approach is based around determining the correctness of the network topology.

In both cases, the existence of a simple Public Key Infrastructure (PKI) is assumed, but a fuller exploration of the key management and key distribution problem has not been undertaken. Other research has explored the possibility of using Identity-based Public Key Cryptography (ID-PKC) [12].

The network topologies are assumed to be unstructured, which places a strong influence on the security relationships which are assumed to exist or be possible.

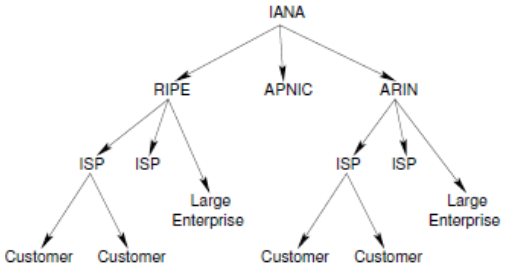

Fig.1. Internet Address Assignment

B. **Inter-domain Routing Security on the Internet**

The network topologies Inter-domain routing on the Internet is managed using BGP4 (Border Gateway Protocol) [19]. This was originally developed for use in a trusted environment, and so provides little security against attackers or misconfiguration. The need for additional security mechanisms has been recognised in recent times, and demonstrated by the AS7007 incident [14][15] and more recent "hijacking" of a part of the YouTube address space [20]. Both incidents are believed to have been due to misconfiguration, rather than malicious intent. The key components in the BGP architecture are Autonomous Systems (as identified by AS numbers), address prefixes, and routers. An Autonomous System is a network domain containing one or more prefixes. Organisations to whom address prefixes and AS numbers are assigned may also be identified as a separate entity. AS numbers and prefixes both form identifier spaces containing globally unique identifiers. The Internet Assigned Numbers Authority (IANA) is the overall manager of these identifier spaces. It delegates parts of these identifier spaces to Regional Internet Registries (RIRs), such as RIPE and ARIN, who then make assignments to end users including Local Internet Registries (LIRs) such as ISPs. This is illustrated in Figure 1. In BGP, networks can be identified as stub or transit domains. This influences the routes they advertise and policies they use. There is also a division between providers and customers. Where customers acquire address space from their upstream provider this prefix is likely to only be advertised by their provider as part of a larger aggregate prefix. Such address space is described as Provider Aggregatable (PA). In other cases end sites may acquire Provider Independent (PI) address space directly from an RIR and advertise this via one or more upstream providers. Current BGP operations depend completely on peers trusting one another not to inject bad information into the routing updates. This is coupled with limited filtering (e.g. to filter out advertisements of unallocated address space, and to ensure that downstream customers only advertise their own address prefixes). In addition to such filtering, there is some use of TCP-MD5 [7] to provide integrity protection for the protocol between peer routers. Recent work has been undertaken in the IETF to understand the requirements for securing BGP. Three general classes of threat against routing protocols have been identified [2]:

- Attacks on the protocol itself
- Falsification of information carried in the protocol
- Forwarding traffic along a different path to that identified by the routing protocol

Threats against the protocol itself include attempts to spoof a network identity for a peer network. These attacks are the simplest to address, since it is essentially a matter of creating a

secure channel between two cooperating entities (neighbouring BGP routers). Falsification is identified as forming a significant set of threats through sending false routing information in the routing protocol:

- Overclaiming: a Byzantine router or an outsider claims ownership of some prefix it does not own, or is not authorized to advertise.
- Misclaiming: an attacker advertises network resources it owns, but in a way different from that intended by the authoritative network administrator. e.g., advertising inappropriate link costs or path lengths.
- Falsification by Forwarders: a router modifies information in the routing protocol that is meant to be immutable.
- Misstatement: an attacker modifies mutable route attributes in an invalid manner. e.g., deleting, inserting or substituting elements of the AS PATH in BGP.

Protecting the routing information against such falsification attacks is a challenging multi-party security problem, and the core focus of research on inter-domain routing security. Verifying that traffic forwarding behaviour is correct requires active monitoring, and again relies on multi-party cooperation to achieve the security goal. There have been a number of different proposals for adding security to BGP, such as S-BGP [11], Secure Origin BGP (so-BGP) [23], and Pretty Secure BGP (psBGP) [22]. These competing proposals, embody different views on the appropriate model for authenticating ownership of identifiers (such as AS numbers and prefixes), even in the relatively well understood Internet case. There are alternative views about whether assertion of AS number ownership should

be provided by the central allocation authorities, or whether Ass should sign certificates to confirm the identities of each other. The same pair of options is also presented for verifying the ownership of address prefixes. These solutions tend to rely heavily on public key signatures, although some attempts are made to ensure that results of signature verification can be cached. Both the computational burden, and the key and certificate storage requirements are significant for a protocol operating on an Internet scale. To address this, other proposals have been made where such signature use is minimised, e.g., Secure path vector (SPV) [9][18]. BGP contains policy components which allow origin networks, transit networks and local policy to affect the choice of path. Because of this, the shortest route is not necessarily the preferred route. Information may be included in route updates to provide hints to other domains on path preferences, and local domains then make a choice from the paths available to them. Policy information published in route update messages needs to be protected against modification. Approaches to BGP security which avoid the use of cryptographic components by relying on BGP policy tools have also been proposed. One solution, pgBGP (Pretty Good BGP) [10], simply adjusts BGP policies to provide some additional cautiousness in accepting new routes. New origin ASs for a prefix are regarded as suspicious for a period of time, and then accepted as normal. This reduces the likelihood of a (shortlived) prefix or sub-prefix hijacking being successful when used in conjunction with appropriate monitoring systems RPSL (Route Policy Specification Language) [1] provides a way for ISPs to describe their routing policies. For example, it will indicate what routes they accept from a particular neighbouring AS, and what routes they advertise to them. This information is stored in one of a number of central databases, and can be automatically extracted to perform filtering on a router. However, deployment is limited and in practice this information tends to be more useful for diagnostics than filtering.

III. TRUST MODEL FOR INTER-DOMAIN MANET ROUTING: SECURITY TRUST MONITOR (STM) DESIGN

We can find in the literature different techniques to secure a wireless Mobile Ad Hoc Network by providing the routing protocols with a more secure implementation ([7], [8], [9], [10] just to reference a few). Protocols like OLSR can be secured by adding cryptography, signatures, timestamps and public key infrastructure (PKI) systems [11]. PKI solutions, among others, also mention the trust concept in the attribution of keys. However, this does not prevent a network of "trusted" nodes from falling victim to a wormhole attack. Like security, trust can be applied at different levels. The security layer we propose provides additional trust confidence based on the discovery of potentially malicious nodes and intruders. The novelty of this approach is to have a middleware security layer that can be used on top of different MANET routing protocols,

given they are provided with a plug-in to the security interface. Using the interface plug-ins, the routing protocols can access the trust values that are defined and calculated at the security layer. The following table shows what action our OLSR plugin will execute for trust values retrieved from the security monitor.

Table I. Trust Values

Trust Value Per Originator Node	OLSR Protocol Action (Input Processing)
0. Normal – Trusted	Accept all packets from originator node
1. Misbehaving – Untrusted	Drop hello messages from originator node
2. Suspicious – Untrusted	Drop hello messages, and do not forward TC messages from originator node
3. Malicious – Untrusted	Drop hello and TC messages from originator node

In Table I an OLSR plug-in takes specific actions based on the trust values it retrieves from STM. The actions can be defined differently and rely on the requirements of the protocol's implementer. Also, if desired, routing protocols can make use of the trust values to implement trust routing [12]. One of the main purposes of our architecture is to populate and evaluate trust metrics outside of the routing protocol layer, thus making it protocol independent. Each protocol defines what action to take based on the trust values retrieved.

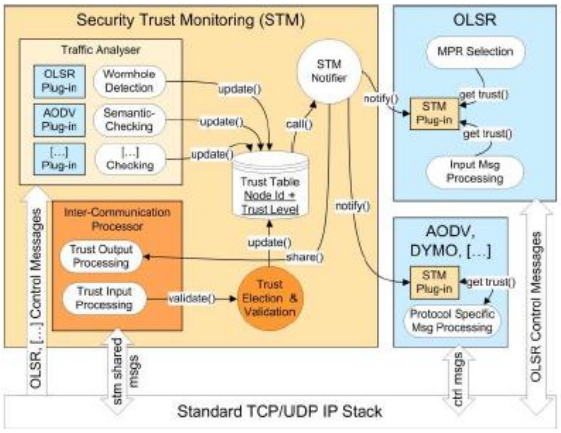

Fig. 2. Overall STM Design OLSR Example

Fig.2. provides an example of the STM architecture adapted to the use of routing protocols such as OLSR.

A. STM's Traffic Analyzer Module

All the traffic going through a node is parsed by the Traffic Analyzer module within the STM and checked against different detection mechanisms (wormhole, semantic, replay, cryptography, etc.). Also, this is where routing protocol plug-ins are defined. These are necessary for STM to parse protocol specific control data and identify abnormal or malicious behavior. I.e. STM is aware of the protocol's packet format used and of any added security encryption, or extra information

fields, such as packet leashes [13]. This module can be updated with new intrusion detection algorithms as they get developed. Depending on the environment where the MANET is deployed and on what hardware, security checking components can be prioritized to minimize impact on the processing performance. The Traffic Analyzer module produces trust values based on locally calculated trust evidence, values that we refer to as local-trust.

B. Inter-Communication of Distributed STMs

An instance of STM is executed in a distributed manner on each MANET node. Every node updates its trust table representation of the whole network based on its own analysis (local-trust). This mechanism is effective against attacks on the network, such as the wormhole. However, there are cases where information about a compromised node needs to be shared among nodes. If a MANET node identifies with certainty (according to predefined parameters) that a node is malicious, STM can send a broadcast message to all nodes to warn them; this is referred as trust sharing. However, as was observed in semantically-based intrusion detection [14], the ability to identify with certainty or very high probability that a node is an intruder when located more than 2 hops away is low. The experiments conducted in that paper had only direct

neighbors conclusively identify attackers. In a nutshell, if the packet judged suspicious originates from more than one hop away, there is little certainty in identifying that the originator node is malicious, because any malicious forwarding node could have modified its content. This is why we consider that trust sharing is only required for each node's immediate (1 hop) neighbors. Restricting the trust sharing to neighbor nodes has the advantage, compared to broadcast flooding, of reducing the impact on the available communication bandwidth. The sharing of trust data among nodes is done within the Inter-Communication Processor module of the STM (see Fig. 1). Periodically, the command to share the trust value of a node is given, and it triggers a message to be generated (Trust Output Processing) and propagated to the neighboring nodes using the standard TCP/UDP stack.

C. STM's Trust Election and Validation Module

Sharing information about discovered malicious nodes, unfortunately, introduces a caveat: STM becomes vulnerable to blackmail attacks. What if a node mistakenly or maliciously sends repeated warnings about a node which is behaving normally? In response to blackmail attacks we introduce 2 counter measures: STM propagating shared messages should use encryption mechanisms and use keys that are different than the ones used within the routing protocol (if any). This is handled at the input and output of STM's message processing. STM should validate the trust values received based on different parameters such as:

o Consistency: is the message recurring?

o Election: is the message shared by other nodes as well? Are a majority of nodes converging to the same trust value? What is the reputation of the node that it is received from?

o Validity: is the value valid?

The validation of the trust values received is handled within the Trust Election & Validation module of STM (see Fig. 1), before calls to update the Trust table are made. The values obtained with evaluation of trust sharing messaging are referred to as global-trust.

D. Trust Table Handling

The values maintained with the trust tables are built out of a combination of local-trust and global-trust calculated values. Local-trust values have a greater weight than global-trust values on the establishment of the table's values since the highest confidence level is accorded to one's self.

E. Return to Trusted State

A node that is once set to be malicious cannot return to its trusted state unless new messages are generated with a "trusted" value. In that case, the validation and election check is applied, and the state change is executed gradually (malicious to suspicious, to misbehaving, to trusted). However, the gradual recovery rule can be overwritten by the owner of a node.

IV. APPROACHES TO SOLVING THE INTER-DOMAIN SECURITY PROBLEM FOR MANETS

An overall architecture for addressing the inter-domain MANET routing security problem is likely to involve a number of components. In particular, cryptographic mechanisms to confirm ownership of identities (e.g., AS numbers) and address space will provide a foundation block. This can then be used to allow the attestation of paths (and related policy information) advertised in the routing protocol. Additionally, verification of packet forwarding behaviour against routing protocol information will assist in identifying misbehavior in the network.

A. Address Ownership

A single approach to the question of address and identity ownership is unlikely to be feasible, since this is strongly driven by the deployment model. Therefore, it is appropriate to consider both hierarchical and peer-to-peer (web-of-trust) models of managing identity and address space ownership.

Real deployments may actually need to use a combination of these at different parts of the allocation hierarchy (e.g., through cross signing of CAs). A standard PKI approach offers a route to providing hierarchical address and network identity assignment. In

the Internet context, this is the approach used by S-BGP [11], which was one of the earlier proposals for securing BGP. It makes use of two hierarchical PKIs using X.509 certificates for assignment of AS numbers and prefixes to organisations. Signatures are used to confirm that a router is authorized to act on behalf of an AS, that an AS is authorized to advertise itself as the origin of a prefix (address attestations). In addition, certificates are provided to routers to confirm that they are authorized to speak on behalf of an AS, and so secure communications between peer routers using IPsec. Certificate distribution is performed out-of-band. However, in the MANET context reliance on a central repository would need to be avoided. Secure Origin BGP (so-BGP) [23] and Pretty Secure BGP (psBGP) [22] mix the use of hierarchical and "web of trust" models for address ownership and AS identities. so-BGP uses a hierarchical X.509 PKI for AS number assignment, with a small number of root CAs, and the possibility of multiple levels in the PKI hierarchy. (The use of multiple levels actually moves this more towards a "web of trust" model, where ASs confirm the ownership of public keys by other ASs). Signatures by ASs are then used to authorize another AS to advertise a particular prefix. This removes the division between the organisation to which a prefix is assigned and the AS that can originate it that exists in S-BGP. psBGP uses a centralised trust model for AS numbers, with a PKI using RIRs acting as certificate authorities to confirm AS bindings to public keys. A decentralised trust model is used to confirm prefix ownership, where a small number of the peers of an AS endorse its ownership of a prefix. In addition, signatures from ASs are used to confirm that BGP speakers (routers) are authorized to act on its behalf. As an alternative to traditional PKI, Identity-based Public Key Cryptography may offer some advantages. In particular, the need to store and exchange certificates is avoided. Some initial proposals have been made on applying ID-PKC to interdomain routing [3]. Where a key management system is in place, the question of revocation also needs to be addressed. This requires the ability to form a revocation decision (either centrally or in a distributed fashion), and often requires a mechanism to distribute this decision to other nodes in the network. Revocation mechanisms are underspecified in most of the current approaches. It is likely to be desirable to make use of the intra/interdomain split in scaling revocation. In particular, it is probably useful to localise individual node revocation to the intradomain context, only propagating it into the inter-domain context in the form of route withdrawals if required. One of the complexities of revocation potentially arises where ownership of a resource (e.g., address space) changes. Re-delegation also adds complexity to the key infrastructure. It may be that, given that changes in ownership are likely to be rare, forcing renumbering in such scenarios is acceptable.

B. Route Correctness

Attacks on routing can target both the control plane (the routing protocol) and the data plane (packet forwarding). The security goal of control plane protocols is to ensure that only valid path updates are propagated on the network. Validity of a path update (e.g.: hA, B, C, D, dsti) is defined according to four primary criteria:

• The router that sent the update was authorized to act on behalf of the domain (A) it claims to represent (by virtue of placing that domain ID in the path).

• The domain from which the update emanates (A) was authorized by the preceding domain (B) in the path to advertise the destination (dst) contained within the update.

• The first domain in the path (D) owns the set of destinations dst (e.g. a set of prefixes), and is thus authorized to advertise them.

• If the update withdraws one or more routes, then the sender must have advertised the route(s) prior to withdrawing it (them).

S-BGP addresses these requirements using the concept of address and route attestation. Address attestation binds a set of destination nodes (dst) to a domain (D). It ensures that a destination (address or prefix) is advertised by a domain if and only if the domain owns the destination. Route attestation created by one domain authorizes a neighbour domain to advertise paths to a set of destinations. Route attestation ensures that a malicious node cannot spoof paths, shrink a path (e.g., create black holes) or alter a path (other than appending itself to the tail of a valid path). so-BGP provides a weaker form of route attestation, providing only assurance of path plausibility. Rather than requiring evidence that B was authorized to advertise the destination, only evidence of the existence of a peering relationship between A and B is needed, which in turn provides the ability to verify the correctness of the AS topology map. In combination with address attestation, this identifies that there is a plausible path to the given prefix.

psBGP uses a model of path consistency checks, where ASs endorse one another's ownership of address space, and peering assertions (i.e. that A is a peer of B) are signed. The combination of the two components allows routes to be verified, based on a seeing a consistent set of assertions from the possible paths available. This approach does, however, provide opportunities for collusion. S-BGP, so-BGP and psBGP rely on a PKI based infrastructures for attestation. Secure path vector (SPV) [9][18] provides an alternative using only efficient symmetric key cryptographic primitives. One of the key drawbacks of these approaches is that the signature size in the update message grows with path length. More recently, Identity-based Sequential Aggregate Signatures (IBSAS) [3] have been proposed

to generate compact signatures that allow signers to attest a path update message as well as the order in which they signed (albeit, incurring higher computation costs than S-BGP and SPV), making it more suitable for a MANET setting. While attestations ensure that all advertised paths are valid, it does not prevent a malicious node from selectively propagating only those advertisements that are of strategic interest (say, with the goal of attracting more paths for traffic analysis or repelling paths to throttle transit traffic, etc). Hence, attestations can only guarantee validity, but not the optimality of path updates. However, if the network is tightly connected (offering multiple alternate paths between any two nodes), then path optimality may not be a serious concern. Techniques which examine the consistency of alternative available paths can also play a role in detecting such modifications. If sufficient alternatives are available then voting could be used to exclude anomalous path options. For example, a route to a prefix received via alternative transits should identify the same origin AS, and potentially the same initial set of components in the AS path. It may also be possible to compare paths to different ASs to identify inconsistencies in the data. In addition, in wireless networks it may also be possible for nodes to monitor the routing messages entering and leaving another router. Such witnesses may be able to identify whether the routes being published by the router are consistent with those it is receiving. This concept is explored more fully in the context of forwarding plane verification later in this paper. There are clearly a number of approaches to verifying the correctness of a route within the routing protocol. In identifying the most appropriate solution, trade-offs will need to considered in terms of protocol overhead, the amount of processing required and the level of guarantee of correctness needed.

C. Verification of Correct Forwarding Behavior

Route correctness ensures that all good nodes commit to a valid path in the control plane. On the data plane, however, nodes may misbehave and forward packets along paths that deviate from their path advertisements. The security goal in the data plane protocols is to ensure that all nodes forward packets in accordance to the paths advertised. This may be achieved through active mechanisms (e.g., probing) or passive mechanisms (e.g., monitoring). One recently proposed active approach, Ordered Multi-Signatures (OMS) [3], allows multiple (ordered) signers to sign the same packet to allow verification of the path it followed through the network. However, OMS is computationally very expensive, making it nearly infeasible to operate at high speeds. One can use secure implicit sampling (SIS) techniques in which: (i) only a small (pseudo-randomly selected) set of packets are signed, and (ii) the selected packets are signed by a (pseudo-randomly selected) sub-set of nodes in the advertised path. Implicit sampling allows the sender (and receiver) to pseudo-randomly determine the subset of nodes in an advertised path that are required to sign a packet. Secure implicit sampling ensures that the choice of nodes that are required to sign a packet is a priori known only to the selected nodes and the sender/receiver. Nonetheless, SIS techniques incur significantly high key management overhead. In a MANET setting one can use cooperative monitoring techniques to verify a node's forwarding behaviour. Cooperative monitoring can be achieved at two levels: neighbor monitoring and end-to-end monitoring. Neighbour monitoring exploits the broadcast nature of wireless channels. A node's neighbour (acting as a witness) can observe packet transmissions by that node. Even if the packet payloads are encrypted, a node can use packet headers to decide if the node is behaving consistently with its advertised path. The radius of the observation area can be increased, so increasing the number of potential observers and the packet forwarding events (and locations) that can be observed. As increasing information is obtained, more detailed determination can be made as to whether "correct" behaviour is being observed. With limited information (i.e., observation only of events nearby), it is possible to determine that neighbouring domains are cooperating, for example, by forwarding packets. By increasing the observation radius, additional packet forwarding events become visible. Consequently, it may then be possible to deduce that rational behaviour is occurring. For example, an observer may then be able to determine that a neighbouring domain is forwarding packets to the neighbouring AS that it said it would. Finally, at the widest level of observation more complex misbehaviours can be observed, such as Byzantine nodes, and that routing is being performed consistently across the network. The determination of observation distance can be performed based on an acceptable risk level. The greater the risk aversion, the higher the degree of observation performed. We remark that a node may exploit MAC-layer misbehavior to appease a witness while ensuring that the transmitted packet does not reach the next node in the network. Fortunately, such denial of service attacks can be easily detected using end-to-end monitoring. End-to-end monitoring techniques rely upon special nodes (monitors) in the network. Monitors may not explicitly participate in the packet forwarding task; however, their goal is to collect end-to-end performance metrics between pairs of monitors. The collected monitoring information is pooled and root cause analysis techniques are employed to detect misbehaving nodes in the network. For instance, let us suppose that two end-to-end tests hM1, A, C, M2i and hM3, B, C, M4i show poor performance (e.g.: heavy packet loss) then the monitors may use this evidence to conclude that C is the most likely suspect. Indeed, the monitors may adaptively run more tests to confirm their hypothesis. However, the monitors have to camouflage their test (or probe) packets amongst normal traffic; otherwise, a malicious router may offer differential treatment to the probe

packets, thereby defeating all probe based end-to-end monitoring techniques. Most monitoring based solutions suffer from the following drawbacks. A subset of monitors may be compromised or malicious and may attempt to actively manipulate the monitored data with at least one of the following two goals: shilling − make a bad node appear good and bad mouthing – make a good node appear bad. In a coalition MANET wherein monitors may belong to different domains whose incentives may not be perfectly aligned, one is faced with the challenge of performing root cause analysis on the monitored data assuming k-out-of-n monitors may be malicious. Trust based mechanisms need to be deployed here to incentivise monitors whose data matches the consensus (e.g.: the monitor's vote matches the majority vote), while punishing those monitors whose data deviate significantly from the consensus. However, one should keep in mind the inherent monitoring error in building such trust based mechanisms. The results of these techniques may allow malicious behavior to be detected, or simply to build up confidence in the correct behaviour of domains under different administrative control. The results of these observations may be used to influence routing policy. For example, some traffic may be restricted to flowing down highly trusted paths, whilst other data may be sent along a partially trusted route, in order that packet forwarding can be observed to build trust.

V. CONCLUSIONS

A security solution to the inter-domain MANET routing problem can be expected to rely on a number of elements offering both cryptographic assurances, and network monitoring Although previous work in intra-MANET routing and interdomain routing on the Internet provide useful background, proposed solutions in those problem spaces cannot be carried over directly to inter-MANET routing security. This paper identifies that key differences arise because of the different organisational structure and more dynamic topologies to be found in military MANET deployments. The issues of address and identifier assignment, and hence their security, are driven by the way the network is managed and deployed. Further, the routing protocol design, its policy components and security are interconnected, and cannot each be considered in isolation. Future work will need to provide strong authentication of the ownership of identifiers, which then supports the ability to verify routes advertised by the routing protocol. Existing research shows the prospect for developing techniques that offer a balance between security and processing or bandwidth overhead. The requirement for these to fit the deployment model will need to be considered, in order to avoid creating a system which is difficult or impossible to use. In addition to identity and path protection, techniques have been identified to provide monitoring functionality which can be used to confirm whether there is misbehaviour on the forwarding path. These can be expected to be used to support verification of correct behaviour by the network.

REFERENCES

[1] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra. Routing Policy Specification Language (RPSL). RFC 2622 (Proposed Standard), June 1999. Updated by RFC 4012.

[2] A. Barbir, S. Murphy, and Y. Yang. Generic Threats to Routing Protocols. RFC 4593 (Informational), October 2006.

[3] Alexandra Boldyreva, Craig Gentry, Adam O'Neill, and Dae Hyun Yum. Ordered multisignatures and identity-based sequential aggregate signatures, with applications to secure routing. In CCS '07: Proceedings of the 14th ACM conference on Computer and communications security, pages 276–285, New York, NY, USA, 2007. ACM.

[4] Chi-Kin Chau, Jon Crowcroft, Kang-Won Lee, and Starsky H.Y. Wong. IDRM: Inter-Domain Routing Protocol for Mobile Ad Hoc Networks. Technical Report UCAM-CL-TR-708, University of Cambridge, Computer Laboratory, January 2008.

[5] Manel Guerrero Zapata. Secure ad hoc on-demand distance vector (saodv) routing, September 2006. Internet Draft draft-guerrero-manetsaodv- 06.txt. ACITA 2008 Page 129

[6] Manel Guerrero Zapata and N. Asokan. Securing Ad hoc Routing Protocols. In Proceedings of the 2002 ACM Workshop on Wireless Security (WiSe 2002), pages 1–10, September 2002.

[7] A. Heffernan. Protection of BGP Sessions via the TCP MD5 Signature Option. RFC 2385 (Proposed Standard), August 1998.

[8] I. W. Ho, B. Ko, M. Zafer, C. Bisdikian, and K. Leung. Cooperative Transmit-Power Estimation in MANETs. In Processings of IEEE WCNC 2008, Las Vegas, March 2008.

[9] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu. SPV: secure path vector routing for securing BGP. ACM SIGCOMM Computer Communications Review, 34(4), October 2004.

[10] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In The 14th IEEE International Conference on Network Protocols, November 2006.

[11] Stephen Kent, Charles Lynn, and Karen Seo. Secure Border Gateway Protocol (Secure-BGP).

IEEE Journal on Selected Areas in Communications, 18(4):582–592, April 2000.

[12] A. Khalili, J. Katz, and W.A. Arbaugh. Toward secure key distribution in truly ad-hoc networks. In Proceedings of 2003 Symposium on Applications and the Internet Workshops, pages 342–346, January 2003.

[13] Y. W. Law, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga. Energyefficient link-layer jamming attacks against wireless sensor network mac protocols. In SASN '05: Proceedings of the 3rd ACM workshop on Security of ad hoc and sensor networks, pages 76–88, New York, NY, USA, 2005. ACM.

[14] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. In SIGCOMM '02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications, pages 3–16, New York, NY, USA, 2002. ACM.

[15] S. A. Misel. Wow, AS7007! http://www.merit.edu/mail.archives/nanog/ 1997-04/msg00340.html, April 1997.

[16] P. Papadimitratos and Z.J. Haas. Secure link state routing for mobile ad hoc networks. In Proceedings of IEEE Workshop on Security and Assurance in Ad Hoc Networks, in 2003 International Symposium on Applications and the Internet Workshops, 2003, pages 379–383, January 2003.

[17] S. Radosavac, J. H. Baras, and I. Koutsopoulos. A framework for MAC misbehavior detection in wireless networks. In WiSE, 2005.

[18] Barath Raghavn, Saurabh Panjwani, and Anton Mityagin. Analysis of the SPV Secure Routing Protocol: Weaknesses and Lessons. ACM SIGCOMM Computer Communication Review, 37(2), April 2007.

[19] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), January 2006.

[20] RIPE NCC. YouTube Hijacking: A RIPE NCC RIS case study. http: //www.ripe.net/news/study-youtube-hijacking.html, February 2008.

[21] M. Srivatsa, B. Ko, A. Beygelzimer, and V. Madduri. Topology Discovery and Link State Detection using Routing Events. under submission.

[22] Tao Wan, Evangelos Kranakis, and P.C. van Oorschot. Pretty Secure BGP (psBGP). In The 12th Annual Network and Distributed System Security Symposium, February 2005.

[23] Russ White. Securing BGP Through Secure Origin BGP. Internet Protocol Journal, 6(3), September 2003.

[24] M. Zafer, B. Ko, and I. W. Ho. Cooperative Transmit-Power Estimation under Wireless Fading. In ACM Mobihoc 2008, Hong Kong, May 2008.