

Securing the Border Gateway Routing Protocol

Bradley R. Smith
Computer Sciences
University of California
Santa Cruz, CA 95064
brad@cse.ucsc.edu

J.J. Garcia-Luna-Aceves
Computer Engineering
University of California
Santa Cruz, CA 95064
jj@cse.ucsc.edu

Abstract

We analyze the security of the BGP routing protocol, and identify a number of vulnerabilities in its design and the corresponding threats. We then present a set of proposed modifications to the protocol which minimize or eliminate the most significant threats. The innovation we introduce is the protection of the second-to-last information contained in the `AS_PATH` attributes by digital signatures, and the use of techniques developed for detecting loops in path-finding protocols to verify the selected route's path information. With these techniques we are able to secure full path information in near constant space, and avoid the recursive protection mechanisms previously assumed necessary.

1 Introduction

Inter-domain routing protocols are designed to perform policy-based routing in an internet of autonomous systems. An autonomous system (AS) is defined as a set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS, and using exterior gateway protocols to route packets to other ASs. In practice, this definition is relaxed to allow multiple intra-domain protocols and several sets of metrics, the focus being on a single administration. Two inter-domain, path-vector routing protocols currently defined are the Border Gateway Protocol (BGP) [20] and the Inter-Domain Routing Protocol (IDRP) [19, 7]; these two protocols are of particular interest because of their current roles as the inter-domain protocols maintaining the global Internet routing infrastructure.

Routing protocols dynamically configure the packet forwarding function in internets which allows for the continued delivery of packets in spite of changes in network topology and usage patterns. These changes typically occur due to the ongoing introduction, failure, and repair of network links and routing nodes, which the protocols have been designed to accommodate. The compromise of the routing function in the global Internet can lead to the denial of network service, the disclosure or modification of sensitive routing information, or, via the reconfiguration of the logical routing structure in the Internet, the diversion of network traffic possibly leading to the disclosure of network traffic to an attacker or the inaccurate accounting of resource usage. Current routing protocols contain few, if any mechanisms to provide for the security of their operation. Those that exist are incompletely defined or are not implemented. Given the evolution of the global Internet to a commercial, production network infrastructure this state of affairs is clearly unacceptable.

⁰This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) under Grant F19628-96-C-0038.

The proposed Internet Security Architecture (ISA) [22] provides an architecture for the inclusion of security facilities in the design of protocols to be used in the Internet. Fundamental to the ISA are four concepts: vulnerabilities, threats, security services, and countermeasures. A **vulnerability** is a weakness in a system's security that may be exploited by an intruder. A **threat** is a potential violation of security, and requires an intruder who has the capability to exploit an existing vulnerability. Threats can be classified into four general categories. *Disclosure* is an event in which an entity gains access to data that the entity is not authorized to receive. *Deception* is an event that results in an authorized entity receiving false data and believing it to be true. *Disruption* is an event that interrupts or prevents the correct operation of system services or functions. And, *usurpation* is an event that results in control of system services or functions by an unauthorized entity.

Vulnerabilities and threats are minimized or eliminated through the provision of six **security services** [17]. *Confidentiality* is the protection of data so it is not made available or disclosed to unauthorized individuals, entities, or processes. *Integrity* is the protection of data so it is not altered or destroyed in an unauthorized manner. *Authenticity* is the verification of the identity claimed by a system entity. *Access Control* is the protection against unauthorized use of system resources. *Non-Repudiation* is the protection against false repudiation of a communication. *Availability* is the assurance that resources are accessible and usable upon demand by an authorized entity.

A **countermeasure** is a mechanism or feature that provides a security service. Examples of countermeasures include encryption of network traffic to provide confidentiality, and the use of challenge-response technology for providing authentication of user logins. The cryptographic tools we will use to implement countermeasures to routing protocol vulnerabilities are primarily encryption and digital signatures. Given these cryptographic tools and the concepts from the ISA, this paper presents a strategy for securing BGP using the following methodology:

1. Analyzing the protocol design to identify vulnerabilities and threats.
2. Identifying the security services needed to reduce or eliminate the vulnerability.
3. Designing the appropriate countermeasures to provide the needed services.

Section 2 states our assumptions, and goals for securing BGP. Section 3 analyzes the security of BGP, and identifies its vulnerabilities and the threats it is susceptible to. Section 4 presents our proposed strategies and countermeasures for securing BGP. Section 5 reviews related work.

2 Assumptions of the BGP Environment

There are four basic components in a BGP system: speakers, peers, links, and border routers [20].

A *BGP speaker* is a host in the network that executes the BGP protocol. *BGP peers* are two BGP speakers that form a connection and engage in a BGP dialog. A BGP peer is either an internal or external

peer, depending on whether it is in the same or a different AS as the reference BGP speaker. The connections between BGP peers are called *links*, with internal and external links being defined similarly to internal and external peers. BGP links are formed using a reliable transport protocol such as TCP. This eliminates the need to implement transport services such as retransmissions, acknowledgments, and sequence numbers in the routing protocol.

A *border router* is a router with an interface to a physical network shared with border routers in other autonomous systems. Similar to BGP speakers, border routers are either internal or external. Note that BGP speakers need not be border routers (or even routers of any kind). It is possible that a non-routing host could serve as the BGP speaker, gathering routing information from internal or other external routing protocols, and advertising that information to internal and neighboring external border routers. This feature is currently in use in the Route Servers of the Routing Arbiter project [5].

We make the following assumptions in designing security mechanisms for BGP:

- The BGP version 4 protocol as defined in [20].
- A BGP speaker can trust its internal peers.
- A BGP speaker can trust information it receives from external speakers only concerning links incident on the AS the external speaker belongs to.
- Intruders have capabilities as described in Section 3.1.
- Key distribution is based on domain names, domain names can be efficiently and securely determined given an IP address of a host, and the key distribution mechanism provides a controllable refresh rate. The DNS Security Extensions [4] might meet these requirements.

3 BGP Threats and Vulnerabilities

We now identify the threats to which BGP is susceptible, and the vulnerabilities these threats exploit. We consider separately threats to the flow of routing traffic and threats to the flow of data traffic that involve portions of the routing infrastructure.

We describe attacks in terms of different classes of internet nodes, including: authorized BGP speakers, authorized BGP routers, and intruders. Authorized BGP speakers are those nodes intended by the authoritative network administrator to perform as a BGP speaker.

3.1 Intruders

We assume that an intruder can be located at any point in the network through which all traffic of interest flows, and that the intruder has the capability to fabricate, replay, monitor, modify, or delete any of this traffic. Interpreting this description for a BGP environment, we identify the following four general classes of intruders: subverted BGP speakers, unauthorized BGP speakers, masquerading BGP speakers, and subverted links.

A subverted BGP speaker occurs when an authorized BGP speaker is caused to violate the BGP protocols, or to inappropriately claim authority for network resources. This typically occurs due to bugs in the BGP software, mistakes in the speaker's configuration, or by causing a BGP speaker to load unauthorized software or configuration information, which can be achieved by many means, depending on the design and configuration of the BGP speaker.

An unauthorized BGP speaker exists when a node that is not authorized as a BGP speaker manages to circumvent any access control mechanisms in place, and establish a BGP link with an authorized BGP speaker. How this is achieved depends on the design and configuration of existing access control mechanisms.

A masquerading BGP speaker occurs when a node successfully forges an authorized BGP speaker's identity. This can be accomplished using the IP spoofing [14] or source routing attacks.

There are a number of forms that a subverted link can take. One is to gain access to the physical medium (e.g. copper or fiber optic cable-plant, the "air-waves", or the electronics used to access them) in a manner that allows some control of the channel. In addition, a link

may be subverted by compromising lower layer protocols in use on the link in a manner that allows control of the channel. An example of such an attack is the TCP session hijacking [8] attack.

3.2 Deception or Disruption of Routing Messages

There are a number of vulnerabilities that allow a strategically placed intruder to fabricate, modify, replay, or delete routing information. With these capabilities, an intruder can compromise the network in a number of ways. The modification or fabrication of routing updates allows an intruder to reconfigure the logical routing structure of an internet, potentially resulting in the denial of network service, the disclosure of network traffic, and the inaccurate accounting of network resource usage. The replay or deletion of routing updates blocks the evolution of subsets of the logical routing structure (in response to topological or policy changes), or resets it to an earlier configuration with results similar to above. Specific attacks include:

- An intruder subverts an authorized BGP speaker.
- An unauthorized BGP speaker establishes a BGP link with an authorized BGP speaker.
- A masquerading BGP speaker takes the role of an authorized BGP speaker in the routing computation.
- An intruder subverts a link through which BGP links pass.

The vulnerabilities these attacks exploit is the lack of access control, authentication, and integrity of BGP message contents.

3.3 Disclosure of Routing Messages

It is relatively easy for an intruder to gain access to routing traffic. The information available from this traffic includes the appropriate next hop to reach a destination, and the path taken by traffic to different destinations. The next hop information is available from other sources, such as monitoring authorized traffic to the desired destination for the next hop it uses, and therefore cannot be protected solely by measures directed at the routing traffic. However, in some circumstances, the path used to reach different destinations may be considered confidential. Specific attacks to obtain this path information include:

- An intruder subverts an authorized BGP speaker.
- An intruder subverts a link through which BGP links pass.

The vulnerabilities these attacks exploit are the lack of confidentiality of peer links, and the level of trust placed in BGP speakers.

3.4 Disclosure of Data Packets

It is relatively easy for an intruder to snoop or disclose data traffic. The vulnerability exploited here is the lack of end-to-end or link encryption services for data traffic. Being beyond the scope of our intended modifications to BGP, we will not address these possible countermeasures further.

3.5 Deception or Disruption of Data Packets

It is relatively easy for an intruder to fabricate, modify, replay, or delete data packets. The effectiveness of these attacks at deceiving or disrupting the source and destination processes depends on the end-to-end protocols in use at the transport layer and above, and is not a routing protocol issue.

However, the effectiveness of these attacks at deceiving the intermediate routing nodes is not an end-to-end protocol issue. Countermeasures to these vulnerabilities will depend on mechanisms in the network or lower layers of the protocol hierarchy. The appropriateness and effectiveness of end-to-end vs. link layer security measures is a fundamental issue in the design of the Internet protocols [11, 21, 23]. While in general these issues do not involve routing protocol mechanisms, two exceptions include the ability to use multiple paths to a single destination, and the inclusion of authentication and access control mechanisms in the packet forwarding function (e.g. [6]). We will not address these measures further in this paper.

4 BGP Security Countermeasures

The general outline of our proposed countermeasures is as follows:

- Encrypt all BGP messages between peers using session keys exchanged at BGP link establishment time. This encryption provides integrity and authenticity of all path attributes whose values are valid for at most one AS hop, and confidentiality of all routing exchanges.
- Add a message sequence number to protect against replayed or deleted messages.
- Add an UPDATE sequence number or timestamp to protect against replayed UPDATE messages.
- Add a PREDECESSOR path attribute indicating the AS prior to the destination AS for the current route. This allows the verification of the path information in a manner similar to that used in path-finding algorithms to detect loops [16].
- Digitally sign all unchanging UPDATE fields whose values are fixed on creation by the BGP speaker originating or most recently aggregating the route. This provides for the integrity and authenticity of not only these fields, but also of the full AS_PATH.

The rest of this section presents a more detailed functional specification of these countermeasures, and an analysis of the effectiveness of these countermeasures against the threats and vulnerabilities identified above.

4.1 Functional Overview of Countermeasures

A number of the following countermeasures are, effectively, implementing secure transport services not available from current transport protocols. Specifically, the peer-to-peer encryption and peer-to-peer sequence number are providing corruption detection, sequencing, acknowledgment, and retransmission mechanisms that are redundant to those provided by TCP. They are required, however, due to the insecurity of the TCP mechanisms [3, 10]. These BGP countermeasures would no longer be required if a secure network [1] or transport protocol [9] were used.

4.1.1 Peer-to-Peer Encryption

Upon establishment of each BGP link, a session key is exchanged by the peers for use in encrypting each BGP message transmitted over that link. One purpose of this encryption is to provide confidentiality of the messages. The other purpose is to provide authenticity and integrity of the KEEPALIVE and NOTIFICATION messages, and of some of the path attributes carried in UPDATE messages.

A number of path attributes carried in UPDATE messages are modified in each AS they transit. These include the NEXT_HOP, MULTI_EXIT_DISC, and LOCAL_PREF attributes. The use of peer-to-peer encryption for authenticity and integrity of these path attributes is based on two points: (a) the recipient of these path attributes receives them from either the most recent modifier or via a single relay that is an internal peer, and (b) our assumption that internal peers are trusted. Given these, peer-to-peer encryption provides a high degree of security in an efficient manner. On detection of corrupted information the link is terminated using a NOTIFICATION message.

4.1.2 Message Sequence Number

A sequence number is added to each message; it is initialized to zero on establishment of a BGP link, and is incremented with each message. On detection of a skipped or repeated sequence number, the BGP link is terminated with a NOTIFICATION message. The size of the sequence number is made large enough to minimize the chance of it cycling back to zero. However, in the event that it does, the link is terminated and a new link is established, resetting the sequence number to zero and establishing a new session key.

4.1.3 UPDATE Sequence Number or Timestamp

A sequence number is added to each UPDATE message to protect against replay. An UPDATE message with a sequence number equal to or less than that of a previously received UPDATE message from the same BGP speaker is defined as invalid and dropped. Given the recommended value of BGP's `MinASOriginationInterval` timer (15 seconds) the sequence number can be relatively small and still be assured of not cycling. Setting this timer to as low as 8 seconds, and assuming a new UPDATE is originated at the end of every interval, a four octet sequence number would last for over 1000 years.

The main difficulty introduced by a sequence number is how to maintain it in the context of arbitrary software and hardware failures. Techniques such as those proposed by Perlman in [18] could be used. However, if cycling of the sequence number must be supported, the following process can be used.

Each BGP speaker maintains an UPDATE message sequence number database on a per BGP speaker <domain name, public key> pair basis. When the cycling of a sequence number approaches, a new public-key pair is generated. The key distribution mechanism and BGP speaker are updated with the new key pair, and the speaker's UPDATE sequence number is reset to zero.

On detection of a change in the public key for an originating speaker, the receiving speaker will add an entry to its UPDATE sequence number database for the new originating speaker <domain name, public key> pair with a sequence number of zero. It will continue to use the old sequence number entry until a sequence number failure occurs where the digital signature validation succeeds using the new entry. At this time the old entry is purged, and the conversion to a new sequence number is complete. Further work is needed on a mechanism to load the database of a newly booted BGP speaker.

Alternatively, a timestamp could be used. The main benefit of a timestamp would be the ease of administration provided by the well-defined external reference for use in resetting lost state. The life of even a small timestamp, while not as dramatic as for sequence numbers, is still significant; assuming a granularity as small as one second, a four octet timestamp still has a life of over 130 years.

4.1.4 Secure AS_PATH Attribute with Predecessor Information

To ensure the authenticity of the AS_PATH attribute we augment UPDATE messages with a PREDECESSOR attribute identifying the AS prior to the destination AS for the current route. We call this AS the predecessor to the destination AS. By including the predecessor information, and a digital signature of this information calculated by the originating router (described in Section 4.1.5), the authenticity and integrity of the complete path reported by a router to any destination can be established by the router's neighbors in a manner similar to that used to detect loops in path finding routing protocols. Specifically, this can be done by means of a path traversal of the verified predecessor information reported by the route.

This information is contained in the new PREDECESSOR path attribute. This path attribute includes the following information: the originating AS, which must be the same as the AS in the AGGREGATOR and the first AS in the first AS_SEQUENCE segment of the AS_PATH path attribute, if these attributes exist; the predecessor AS which must be the same as the second AS in the first AS_SEQUENCE of the AS_PATH attribute, if it exists; the IP address of the originating BGP speaker, which must be the same as the IP address in the AGGREGATOR attribute, if it exists; and a TYPE field which can take on the value of either ADD or DELETE.

The ADD version of the PREDECESSOR attribute is generated by the speaker that originates the UPDATE message. This may either be the creator of an unaggregated UPDATE, or the last speaker to perform an aggregation of the routing information in the current UPDATE. The purpose of this form of the attribute is to identify the originating BGP speaker whose key is used to digitally sign the UPDATE, and to identify the destination and predecessor information in the absence of AGGREGATOR and AS_PATH attributes (see below regarding transit-only UPDATES).

The DELETE version of the PREDECESSOR path attribute serves the purpose of identifying a previously reported predecessor relationship that is no longer valid. Possible reasons for this change include the failure of an inter-AS link, or the termination of a transit traffic agreement. This segment type may be generated by either end of the deleted link; the originating AS field of the PREDECESSOR attribute specifies which BGP peer this was.

Due to the policy-based selection and propagation of routes in BGP, it is possible that an AS could be used on a path to a destination while it is not reachable as a destination itself. To ensure propagation of predecessor information for such transit-only ASs to all potential source ASs of transit traffic, an UPDATE message can be sent with a PREDECESSOR attribute, the minimal required set of other attributes, and no NLRI information.

To detect when such transit-only predecessor information should be transmitted, each BGP speaker must track what predecessor information has been forwarded to each neighbor. When a route is selected for propagation to a neighbor, any predecessor information implied by the route not already transmitted to the neighbor must be sent before the route itself is sent. How to handle any lack of predecessor information by either the sender or receiver of an UPDATE is a policy decision.

This predecessor information is used by each node to, conceptually, maintain a *predecessor table* similar to the routing table used in path-finding algorithms. Specifically, the predecessor table is a column vector containing the distance of the chosen path to each destination, and its corresponding predecessor and successor information. It is updated as each PREDECESSOR attribute is processed.

The information maintained in the predecessor table is used to verify an AS_PATH attribute. Before a speaker selects a route for use, that route's AS_PATH attribute should be verified by a walk through the predecessor table. The timing of this verification is not specified, being influenced by the expected frequency of invalid AS_PATH attributes, expected load, and the performance requirements of the speaker. Options for when to perform this verification include on receipt of the AS_PATH, or on selection of the route for use.

4.1.5 UPDATE Message Digital Signature

To ensure the integrity and authenticity of the unchanging UPDATE message information, it is digitally-signed by the originating BGP speaker specified in the PREDECESSOR attribute. Without protection, trust of this information requires trust of BGP peers regarding information concerning links not incident on their AS. This is something we explicitly do not do. By including the PREDECESSOR attribute information in this signature we protect, in addition to the information in the current UPDATE, the full path information contained in the predecessor table described above.

The UPDATE message digital signature is stored in the Marker field of the header, and is calculated over the following fields: UPDATE sequence number, Unfeasible Route Length, Withdrawn Routes, ORIGIN, ATOMIC_AGGREGATE, AGGREGATOR, PREDECESSOR, and the NLRI. This definition of the digital signature assumes that these fields are only meaningful as a unit; that a change in one requires the re-computation of them all. If the protocol evolves to where this is not the case, and subsets of these attributes may be updated independently by different BGP speakers, additional sequence numbers and associated digital signatures will be introduced.

Figure 1 illustrates the proposed modifications using the UPDATE message, which includes all proposed new fields, as a model.

4.2 Countermeasure Effectiveness

Referring back to Section 3, we now analyze the impact of each countermeasure on the identified threats.

Deception of Routing Messages: The digital signature protects against the fabrication and modification of Withdrawn Routes, ORIGIN, AS_PATH, ATOMIC_AGGREGATE, AGGREGATOR, and NLRI information in the UPDATE message by subverted speakers. Peer-to-peer encryption protects against the fabrication or modification of BGP messages by subverted links.

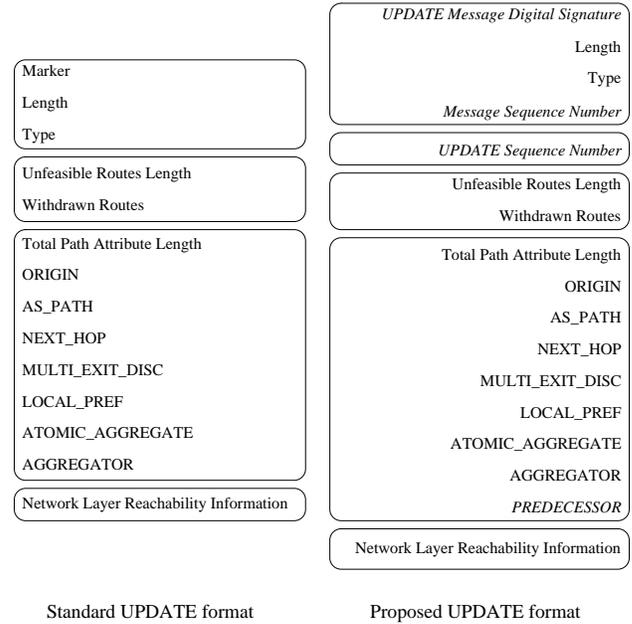


Figure 1: Proposed UPDATE Message Changes

Disruption of Routing Messages: The message sequence number protects against the replay or deletion of BGP messages by subverted links. The UPDATE message sequence number protects against the replay of the UPDATE message information listed above by subverted speakers.

Disclosure of Routing Messages: The encryption of BGP messages protects against the disclosure of routing messages by subverted links.

Threats and Vulnerabilities not Addressed:
 The deletion of UPDATE messages by authorized BGP speakers.
 The disclosure of routing messages by authorized BGP speakers.
 These vulnerabilities are inherent in the need to trust BGP speakers to accurately represent their routing policies, and to maintain confidentiality of routing information.
 Provision of access control to protect against masquerading BGP speakers is still dependent on implementation.

4.3 Performance Analysis

The cost of these countermeasures is in the space for the new sequence numbers and digital signatures, and the time for computing encryption and digital signatures, and verifying these protections. From the perspective of the actions occurring in a BGP system, the costs are as follows:

Message generation and reception: Space: A new field is added for the peer-to-peer sequence number. **Time:** The cost of a symmetric key encryption and decryption of each message.

Initiation and reception of UPDATE messages: Space: The Marker field is used for the UPDATE message digital-signature. Each UPDATE message includes a new UPDATE sequence number and PREDECESSOR attribute. **Time:** The time to perform the computation and verification of the UPDATE message signature.

Route Selection: Time: The time to verify signatures for each link. This cost will only be incurred twice for each used link: once for the ADD and once for the DELETE.

5 Related Work

Kumar [12] analyzes the security requirements of network routing protocols, and discusses the general measures needed to secure the distance-vector and link-state routing protocol classes. He identifies two sources of attacks: subverted routers, and subverted links. Since attacks by subverted routers are seen as difficult to detect, and of limited value to the intruder, Kumar focuses his attention on securing protocols from attacks by subverted links. For distance-vector protocols this translates to the modification or replay of routing updates. The specific countermeasures proposed by Kumar are neighbor-to-neighbor digital signature of routing updates, the addition of sequence numbers and timestamps to the updates, and the addition of acknowledgments and retransmissions of routing updates. Kumar and Crowcroft [13] perform a similar analysis of inter-domain protocols, and come to similar conclusions as the previous paper for providing security of distance-vector related routing protocols (they specifically address the path-vector routing protocol IDRP). The one addition they make is to encrypt neighbor-to-neighbor updates.

These results are similar to ours with the exception that we explicitly assume the existence of subverted routers, and provide countermeasures to protect against them. We feel this is important as BGP speakers are potentially vulnerable to attacks from a number of sources, with potentially catastrophic results from success.

Murphy [15] outlines a solution for securing distance-vector protocols that involves including the information used to select a route, signed by the neighbor it received it from, in the routing update it then signs and transmits to its neighbors. She points out that this requires the validation of a number of nested signatures equal to the number of routers in the path. This results in both update size and validation computation time problems as the size of the network grows. These problems result, fundamentally, from the redundant signing of link information for paths that are supersets of paths used to reach destinations traversed in the longer path. We avoid these problems by signing only the component link information, in the form of predecessors, and performing a path traversal to validate full paths. This results in the use of constant space, and significantly reduced computation time.

6 Concluding Remarks

In this paper we analyze the security weaknesses of the BGP protocol, and identify a number of threats involving the deception, disruption, and disclosure of routing message traffic. We propose countermeasures that eliminate or minimize most of these threats. The primary innovation we introduce is the protection of the predecessor information contained in the `AS_PATH` attribute with a digital signature which is used to verify full paths using techniques used in path-finding protocols for detecting loops. Using these techniques we are able to secure full path information in constant space, and avoid the recursive protection mechanisms previously assumed necessary.

In summary, we show that it is possible to effectively and efficiently secure the BGP routing protocol. Our primary means to accomplish this is the cryptographic protection of the predecessor information existing in the BGP protocol, and techniques developed in path-finding routing protocols [2, 16].

References

- [1] R. Atkinson. Security Architecture for the Internet Protocol. RFC 1825, Aug. 1995.
- [2] C. Cheng, R. Reley, S. P. R. Kumar, and J. Garcia-Luna-Aceves. A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect. *Computer Communications Review*, 19(4):224–336, 1989.
- [3] W. Diffie. Security for the DoD Transmission Control Protocol. In *CRYPTO '85*, pages 108–127. Springer-Verlag, 1985.
- [4] D. E. Eastlake 3rd and C. W. Kaufman. Domain Name System Security Extensions. Internet draft: draft-ietf-dnssec-secext-10.txt, Aug. 1996.

- [5] D. Estrin, J. Postel, and Y. Rekhter. Routing Arbiter Architecture. <ftp://ftp.isi.edu/pub/hpcc-papers/ra/ra-arch.ps>, June 1994.
- [6] D. Estrin, M. Steenstrup, and G. Tsudik. A Protocol for Route Establishment and Packet Forwarding Across Multidomain Internets. *IEEE Trans. on Networking*, 1(1):56–70, Feb. 1993.
- [7] International Standards Organization. *ISO/IEC 10747: Information technology - Telecommunications and information exchange between system - Protocol for exchange of inter-domain routing information among intermediate systems to support forwarding of ISO 8473 PDUs*, Aug. 1994.
- [8] L. Joncheray. A Simple Active Attack Against TCP. In *Proc. 5th UNIX Security Symposium*, pages 7–19. The USENIX Association, June 1995.
- [9] L. Joncheray. Public Key Encryption Support for TCP. Internet Draft: draft-joncheray-encryption-00.txt, May 1995.
- [10] S. Kent. Some Thoughts on TCP and Communication Security. MIT Laboratory for Computer Science, Local Network Note No. 6, Apr. 1977.
- [11] S. Kent. Comments on “Security Problems in the TCP/IP Protocol Suite”. *ACM Computer Commun. Review*, 19(3):10–19, July 1989.
- [12] B. Kumar. Integration of Security in Network Routing Protocols. *ACM SIGSAC Review*, 11(2):18–25, Spring 1993.
- [13] B. Kumar and J. Crowcroft. Integrating Security in Inter-Domain Routing Protocols. *ACM Computer Commun. Review*, pages 36–51, 1993.
- [14] R. T. Morris. A Weakness in the 4.2BSD Unix TCP/IP Software. Technical Report 117, AT&T Bell Laboratories, Murray Hill, New Jersey 07974, Feb. 1985.
- [15] S. L. Murphy. Presentation in Panel on “Security Architecture for the Internet Infrastructure”. Symposium on Network and Distributed System Security, Apr. 1995.
- [16] S. Murthy and J. Garcia-Luna-Aceves. An Efficient Routing Protocol for Wireless Networks. *ACM MONET Journal*, 1996. Special issue on Routing in Mobile Communication Networks.
- [17] D. Nessel. The Internet Security Architecture. In *Proceedings: Internet Security Workshop*. IEEE, Nov. 1994.
- [18] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. Report MIT/LCS/TR 429, Massachusetts Institute of Technology, Oct. 1988.
- [19] Y. Rekhter. Inter-domain Routing Protocol (IDRP). *Internet-working: Research and Experience*, 4:61–80, 1993.
- [20] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771, Mar. 1995.
- [21] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-End Arguments in System Design. *ACM Transactions on Computer Systems*, 2(4):277–288, Nov. 1984.
- [22] R. W. Shirey. Security Architecture for Internet Protocols. A Guide for Protocol Designs and Standards. Internet Draft: draft-irtf-psrg-secarch-sect1-00.txt, Nov. 1994.
- [23] V. L. Voydock and S. T. Kent. Security Mechanisms in High Level Network Protocols. *ACM Computing Surveys*, 15(2):135–171, June 1983.