Building Secure High Speed Extranets

Chandana Gamage Jussipekka Leiwo Yuliang Zheng

Peninsula School of Computing and Information Technology
Monash University
McMahons Road, Frankston, Vic 3199, AUSTRALIA
Phone +61-(0)3-9904 4287, Fax +61-(0)3-9904 4124
E-mail:{chandag,skylark,yuliang}@fcit.monash.edu.au

Abstract

Extranets are a technology for creating logical views of geographically separate LANs by providing a transparent interconnection mechanism between them over a WAN. As LAN bandwidths are typically higher than those of a WAN, and LANs typically support stronger security features, it is essential that extranets can be constructed to be both secure and high of speed. The security of an extranet can be significantly reduced if the interconnection mechanism is not capable of providing comparable level of security to that of LANs being interconnected. Typically, high speed and security require trade offs, but in this paper a method shall be proposed for employing IP switching technology and a new public key cryptographic paradigm, digital signcryption, to construct extranets with both high speed and high level of security.

1 Introduction

An extranet is a private virtual network that interconnects several local subnets (or intranets) over a wide area network (WAN) to carry IP datagram [13] traffic. A typical example of an extranet is the interconnection of geographically distributed local subnets belonging to a single organization using the global Internet. The individual subnets that are part of an extranet are usually local area networks (LAN) with completely localized management and control over its operation. Another type of extranet can be constructed by interconnecting parts of intranets belonging to several organizations through the public Internet to form a shared private internet.

The main identifying attribute of an extranet that differentiates it from a LAN or a WAN is its administrative control structure. A LAN operates under centralized local control policies while a WAN operates under distributed cooperative control. In contrast, an extranet is a collection of *logical views* of private intranets (LANs) that are interconnected over the public Internet to form a private virtual network. For example, this logical view can a be portion of a local subnet partitioned through a packet filtering gateway or permission to execute only a controlled set of applications by blocking external access at host network connectivity ports. Unlike in normal LAN or WAN implementation, security is an integral design feature of extranets. Additional design features are interoperability, reliability and performance.

Interoperability concerns are addressed through standard Internet protocols and related technology on which extranets are constructed of. If the underlying network infrastructure is

solely based on IP technology, network layer interoperability is not a significant issue. However, given the heterogeneous nature of the Internet, achieving of application layer interoperability requires significant additional work. Typical approach is deployment of middleware such as CORBA/ORB and DCE to support distributed applications. On application layer, extranet provides participating organizations with a network infrastructure to integrate their interacting business processes by executing distributed applications and sharing data. Such a distributed on-line systems allow implementation of many of the models for carrying out electronic commerce in a secure, efficient and reliable manner. Example applications for extranets include electronic data interchange (EDI) based systems such as just-in-time (JIT) manufacturing and inventory control [5], work flow document management and real time funds and stock portfolio management. High performance extranets can be used for collaborative interactive work such as industrial design and modeling or real-time multimedia conferencing.

Reliability of extranets at the network layer is based on the same robust and fault-tolerant IP datagram based network model on which Internet operates at. Provision of the necessary performance for high bandwidth and delay sensitive applications together with secure operation of extranets are the areas in which new solutions are required, they have not been traditionally treated as core issues by Internet developers. Therefore, it is justified and important to study them from an extranet construction point of view. Within this paper they are studied in concert with the aim of providing a secure-integrated research approach towards construction of high speed extranets.

The paper is started by surveying models for building extranets in section 2. This survey is then followed by an outline of methods for constructing high speed extranets in section 3. Section 4 provides with an overview of an integrated IP switch firewall. Conclusions are drawn and directions highlighted for future work in section 5.

2 Security in Extranets

Extranets for low bandwidth applications can use available Internet connections for routing of IP datagrams between participating intranets over the public WAN as shown in figure 1. For extranets supporting applications with sustained high bandwidth requirements, leased communication circuits may be necessary. The connectionless datagram based IP model is a flexible and robust networking mechanism over failure prone WANs. However, in an extranet environment, it is necessary to support integrated services with diverse quality of service (QoS) requirements for bandwidth allocation, delay and throughput over this network infrastructure. For this purpose, IP model can be extended by adding a capability to cache soft states relating to the flow of packets across internetworking units such as routers and switches.

An example of this method is the reservation protocol (RSVP) that does bandwidth allocation at the network layer to increase the throughput by reducing the potential for packet loss due to buffer overflow and subsequent retransmissions [15]. The RSVP protocol mechanism is limited in its capacity to make significant improvements to the end-to-end latencies as no attempt is made to speed up the processing of individual network protocol data units (NPDUs) as they pass through an interworking unit. As the focus of this paper is on high speed extranets including high packet throughput at routers, section 3 presents a more detailed discussion on this aspect.

In extranets, security gateways or firewalls provide secure external network access to trusted hosts located internal to a local subnetwork. Multiple security gateways operating in coordi-

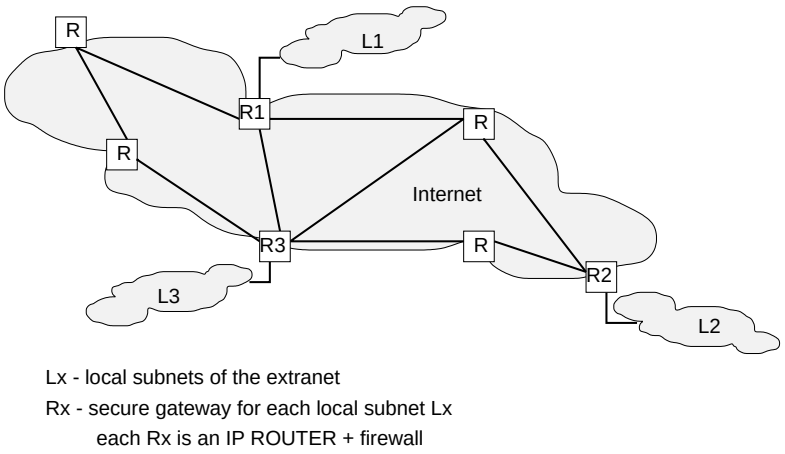

Figure 1: Building extranets using firewalls

nation are used to build the secure virtual private network for the extranet. A gateway based secure networking environment can efficiently implement a host-oriented keying scheme (as against a user-oriented keying scheme) by performing all the *network layer* security related processing at the gateway thus limiting the number of hosts involved in the key management scheme. In this mode of operation, the secure gateway functions as a key management proxy for the trusted hosts inside its subnet.

Enforcement of security can be divided into two categories: provision of security at the point of entry to a network and provision of security of transmitted data. Security at the point of entry is usually enforced by authentication and access control schemes and protection of data during transmission by authenticity, confidentiality and integrity services.

2.1 Security at the Point of Entry

A firewall is a single point of entry to a protected intranet and provide its first line of defense against *outside* attacks. However, firewalls are not transparent in their operation due to the need for close interaction with supported applications. A firewall can block or grant access based on a combination of criteria including source and destination IP addresses (by IP packet level filtering), network port connection type (by TCP/UDP frame level filtering) and application specific attributes such as user authentication credentials (by application frame level filtering) [11]. The cost of filtering related processing increases as the protocol stack layer at which we perform the function increases. However, the flexibility and range of filtering that can be carried out also increases as we go up the protocol stack giving the extranet designers with a performance/flexibility trade-off. In actual implementations, an extranet built using firewalls is likely to perform filtering at each level of the protocol stack. The main disadvantage of firewall technology is that it is modeled on perimeter security from external attacks, therefore limiting its protection capabilities against internal attacks.

Access control systems play a major role in secure extranets. While a firewall is tasked with prevention of unauthorized entry to an intranet at a much coarser granularity, an access control policy and mechanisms with associated authentication techniques are used to extend the access authorization at a much finer granularity to individual hosts, databases or application servers within the intranet.

2.2 Security of Transmitted Data

Apart from its role as a protection boundary, a firewall operating at the network layer is capable of providing host-to-host security. The IETF proposal for a security architecture for the Internet [3] at the IP layer (IPSEC) can be used to provide data origin authentication, confidentiality and integrity for connectionless IP datagram delivery. IPSEC defines two security specific headers: IP authentication header (AH) and IP encapsulating security payload (ESP) header. AH is designed to provide only integrity and authentication for IP datagrams [1] while ESP is designed to provide confidentiality along with integrity and authenticity for datagrams [2] transfered through a network connection. ESP can operate either in *tunnel- mode* by encrypting the complete IP datagram and appending a new clear-text header or in *transport-mode* by encrypting only the upper layer PDU contained in the datagram payload.

For secure communication in IPSEC, end-point nodes must establish a security association (SA) [3]. An SA which include security parameters relevant to a particular network connection is uniquely identified by the combination of an IP destination address and a security parameter index (SPI). An SA for IPSEC is generally one-way and contains following information for use by the receiver:

1. Authentication algorithms, their modes and cryptographic keys used for authentication.
2. Encryption algorithms, their modes and cryptographic keys used for encryption. Also, if an initialization vector (IV) is used, its size.
3. Cryptographic key lifetimes or the key update event trigger.
4. The lifetime of the SA and its IP source address.
5. Security label for the protected data. This may conform with a label hierarchy as used in multi–level secure systems.

IPSEC support interoperability by adopting cryptographic algorithms widely used by the Internet community as the default value set for SAs. However, this default cryptographic algorithm set is based on symmetric key cryptosystems (such as MD5 and DES) and does not provide non–repudiation of IP datagram transmission. IPSEC decouples key management mechanism from the task of secure transmission of datagrams. This separation of two streams of activity allow for greater implementation flexibility.

Firewalls need to process IP datagram headers to obtain packet specific information such as source and destination IP addresses. They also need to process datagram payloads to obtain upper layer protocol specific information such as port numbers and protocol types. Such information is used by the firewall to perform appropriate packet filtering and access control. Use of AH in IPSEC in any mode does not affect the operation of firewalls. However, firewalls will not be able to operate correctly if ESP is used with host-oriented keying without firewall access to the corresponding SA [3].

We assume the existence of a supporting infrastructure for successful implementation of a secure extranet. Major elements of such an infrastructure would be digital public key certificate management facilities, trusted third party key servers and meta directory services for obtaining information about participating intranets, hosts, users and applications.

3 Building a High Speed Extranet

To obtain higher speeds of transmission, we need to use high bandwidth network connections when constructing extranets. Two of the main techniques available are

1. Use of high speed dedicated lines for site interconnection.

   Building an extranet using high bandwidth leased lines to create a fully connected mesh over the WAN is an expensive option as well as decreasing the reliability of the network infrastructure. This method would loose the main advantages of using IP datagram routing as the network transmission technology, low cost and fault resilience, in which routing software chooses an appropriate low cost path for hop-by-hop forwarding of IP datagrams while routing around failed or congested links. Therefore, it shall not be further considered herein.

2. Use of high speed routers to connect to a high bandwidth backbone network.

   This method is feasible as major portions of the Internet are supported by high bandwidth optical transmission links provided by major telecommunication carriers. It also has the additional advantage of preserving favorable properties of IP datagram routing, low cost and error resilience. The remaining major requirement for a successful extranet implementation under this option is a high speed routing device.

The addressing model of the Internet and the format of the IP layer PDU results in considerably low performance for the two main tasks carried out by routers:

Packet forwarding: routing table lookup is a sequential search for a longest matching prefix over a relatively large address space.

Packet copying: as IP datagrams are variable size units, copying from an input port to an output port of a router involves computationally intensive buffer management activity.

In comparison, ATM switches perform the corresponding tasks with a much higher throughput:

Cell forwarding: virtual circuit (VC) table lookup is a direct indexed access.

Cell copying: ATM cells are small fixed size units facilitating fast hardware switching between ports and simpler buffer management.

The common approach to building high speed routers has been to integrate link layer (layer 2) switching with network layer (layer 3) routing. The various schemes based on this technique include IP switching [9, 10], Tag switching [14], IP over ATM [12] and CSR [4]. The higher throughput of such a hybrid router is achieved by efficient cut-through switching of fixed-size cells at layer 2 to bypass the slow routing of variable-size packets at layer 3.

The rest of this section is dedicated in the research on IP switching and security. First, an overview is provided of IP switching and then the implications of IP switching to the network layer security are discussed.

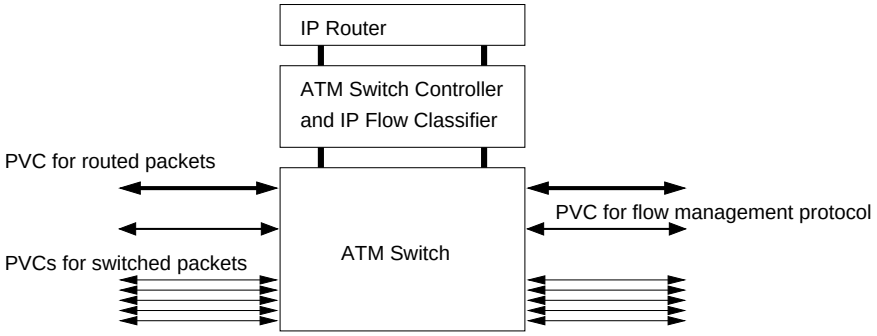

Figure 2: Block structure of the IP switch

3.1 Overview of IP Switching

The operation of an IP switch is based on the concept of a sequence of IP datagrams, termed a *flow*, characterized by a set of common attributes such as source and destination IP addresses, protocol type, port number, etc. The first few datagrams belonging to a flow are routed by the IP switches as in normal datagram routing (*full routing*) at layer 3 and the flow management software in IP switches profiles this flow of datagrams to decide on its suitability for direct switching (*cell streaming*) at layer 2. In Ipsilon IP switches this process is handled by the general switch management protocol (GSMP) [7] and Ipsilon flow management protocol (IFMP) [8]. The decision to build an end–to–end ATM VC for a particular flow is first made by IP switches towards the destination address (downstream) and the connection establishment gradually propagates towards the source address (upstream). Ipsilon's IP switches does not use standard ATM switching protocols and software that conform to ATM–Forum or ITU–T standards and instead use GSMP for switch control and IFMP for link–by–link VC creation and management.

In the cell streaming mode, per-packet processing overhead of full routing is reduced through simple label swapping using connection state tables maintained at each IP switch. This is a compromise between the state–less packet routing in which a routing decision is made for each packet and the state–full (or hard state) packet switching in which an end-to-end connection is pre-established for a complete packet flow. The routed packets are transmitted between IP switches using a default ATM permanent VC (PVC) after encapsulating with a logical link layer/subnetwork attachment point (LLC/SNAP) header for an ATM adaptation layer–type 5 (AAL5) frame. The switched flows are assigned a PVC from a pre-established pool of VCs that each IP switch maintains with its adjacent IFMP–compliant peers (see figure 2). All routed packets are reassembled at the router for routing and flow classification decision making. The IFMP redirect messages sent by a downstream switch to an upstream switch initiates the VC for the switched link for direct hardware switching of the IP datagrams after encapsulating to an AAL5 frame and segmented in to ATM cells.

As shown in figure 3, each local subnet uses an IP switch as both the secure gateway and the border router to connect to the Internet. The maximum performance advantage of IP switching can be obtained only if there is an end-to-end network path through intermediate IP switches. Otherwise, only partial speed increases can be gained for datagram delivery with some segments of a path being switched at layer 2, while remaining segments are routed at layer 3. This scenario is illustrated by the routes indicated in figure 3 showing multiple redundant paths between the subnetworks belonging to the extranet. While each subnet can make an end–to–end connection

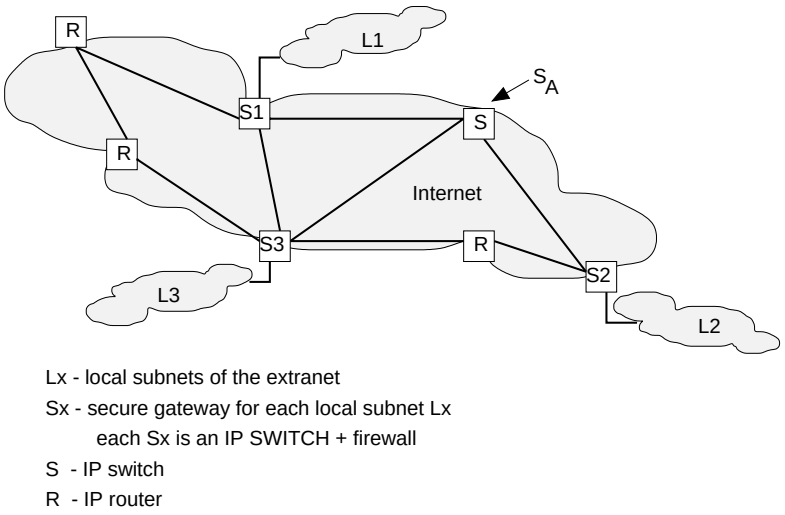

Figure 3: Building high speed extranets using IP switches

through IP switches, if the IP switch S_A were to fail, all datagram traffic to and from subnet L_2 will have to be routed at IP switch S_3. The expected high performance of IP switches can be fully achieved only in a network of IP switches that interconnect with each other to run the IFMP protocol to provide end-to-end layer 2 bypass for packet flows. A simulation study in [6] shows an appreciable performance improvement for a network of adjacent IP switches with a high percentage of datagrams being switched in many of the simulated environments.

3.2 IP Switching and Network Layer Security

If the screening router of a firewall is implemented using an IP switch, it will be able to filter only those packets that are processed under full routing mode. Once the flow of packets is switched over to cell streaming mode, the router no longer has access to the packets for filtering. Therefore, when we bypass (or cut-through) layer 3 processing to achieve higher speed, some of the firewall safeguards will also be bypassed. However, in the extranet model shown in figure 3, IP switches are used as the edge routers of the subnets and receive all outgoing packets and reconstructs all incoming packets. Therefore, the IP switch can perform all firewall related packet filtering based on the IP and TCP header information. Furthermore they can execute IPSEC authentication and encryption tasks on the received packets (after reassembly) and transmitted packets (before segmentation).

When IP datagrams are switched on an ATM VC, the header fields used for flow classification (IP addresses, protocol type, etc) are removed and a compressed header is used for encapsulating the datagram to a AAL5 frame. The removed fields are stored and associated with the corresponding VC for use in header reconstruction at the time reassembling the cells into a packet. This is done to provide an additional measure of security so as to prevent an attacker from establishing a connection to an allowed port through a firewall and then changing the header fields of packets sent through that VC tunnel to gain access to unauthorized hosts or ports [10].

4 An Integrated IP Switch Firewall

In section 2 we have discussed the construction of extranets using standard building blocks (such as Firewalls and IPSEC) available for use in the Internet. In section 3 we discussed improving the network throughput of extranets by replacing standard packet routers with hybrid IP switches. Although switching can increase network bandwidth, continuing use of IPSEC by the IP switch based approach does not address improvements in the efficiency of cryptographic processing needed for secure high speed communication. In this section we propose a modification to the operation of an IP switch to improve its efficiency with respect to security processing.

For secure IP datagram transmission between edge routers functioning as security gateways for extranets, IPSEC must be used in tunnel mode where datagram transmission over the public WAN is always between the gateways through secure virtual tunnels. This allows a gateway to act as a secure proxy for the trusted hosts inside the local subnet and locally deliver or receive datagrams to or from the hosts. Even with secure tunnels, secure datagram transfer and processing is still done on a per–packet basis as the IP model of communication is stateless. However, once a flow is identified for a sequence of datagrams, IP switches begin to maintain soft–state for those switched packet transfers. This gives us an opportunity to increase the speed of security processing by amortizing the cryptographic computational costs over a longer sequence of datagrams rather than on a single datagram as is the case in stateless IPSEC. We achieve this by providing security functionality for the packet stream at the ATM cell level rather than at IP packet level.

It is important to note that the above suggested scheme can be implemented only if the switched flow is established on an end–to–end basis from the source gateway to destination gateway giving a continuous soft–state for the flow. If packet routing is done in intermediate segments of a link, then it is not feasible to map the ATM cell level security association into the IPSEC datagram level security association as intermediate IP switches cannot perform authentication and encryption on behalf of end–point gateways to reconstruct the packets. Above scheme requires several additions to the existing switch and flow management protocols. We briefly outline the main points below

Establishing an ATM cell level security association

1. When a *source gateway* IP switch receives an IFMP redirect message from its downstream neighbor, it will establish the switched flow as per the standard protocol and additionally send a new IFMP *probe request message* to determine if the established flow is end–to–end.

2. If an intermediate IP switch receives a probe request message for a flow that it had not switched onto one of its downstream neighbors, it will send a *negative probe reply message* to the upstream switch. Otherwise, it will forward the probe request onwards to the appropriate IP switch.

3. If an intermediate IP switch receives a probe reply message from a downstream IP switch, it will forward it to the upstream IP switch of the associated flow. If the intermediate IP switch has terminated the upstream switched flow, it will simply discard the received probe message.

4. When a *destination gateway* IP switch receives a probe request for a switched flow, it will send a *positive probe reply message* to the upstream switch.

Using a suitable time–out value and a retry policy, above sequence of steps will enable a source gateway to determine the existence of an end–to–end switched flow. The payload of the probe request message can be used to transmit the set of SA parameters for the associated flow. Conversely, the payload of the probe reply message can provide an acknowledgment for the received SA.

Flow encryption key and secure key transport

Apart from the SA, the payload of the probe request message can securely transport a cryptographic key. This flow encryption key (FEK) is for the encryption of IP datagram payloads to provide confidentiality while being transmitted over the switched link. Also, the FEK in combination with a suitable keyed hash function can also be used to generate a message authentication code (MAC). The MAC can be generated either on a per–datagram or a per–flow basis. For applications such as Telnet in which packets are used interactively, each datagram needs to carry a MAC to ensure its authenticity and integrity. For applications such as FTP, an entire flow can have a single MAC to adequately secure the transmitted content. This MAC generation policy is influenced by application layer considerations and can be parameterized into the SA.

Another important aspect is the secure authenticated transportation of the FEK from the source gateway to the destination gateway. IP switch control messages used by the Ipsilon designed protocols are formatted to fit into single ATM cells for efficient transmission and processing. Secure transmission of a symmetric key of reasonable length (e.g. 70 bits) in a single ATM cell utilizing traditional public key crypto systems such as RSA and ElGamal is not possible due to the very long cryptogram (more than 512 bits) generated by those algorithms in the sign and encryption process. However, the new public key primitive, Signcryption [16], can generate a cryptogram that can fit into the limited payload (maximum of 384 bits) of an ATM cell. Therefore, the flow management protocol can be augmented to efficiently carry a FEK in a single protocol message.

Operation of the extended security protocol

We use IPSEC to provide secure communication when the datagrams are routed or switched by IP switches prior to the establishment of an end–to–end switched flow using probe messages. The security related processing under this mode is external to the IP switch and functions separately in the firewall. However, once an end–to–end SA is established, the security processing needs to be closely integrated with the switch and flow management protocols. For instance, if an intermediate IP switch discontinuous the end–to–end flow by issuing an IFMP reclaim message and reverting to packet routing, then the source gateway must be notified to change over to IPSEC style security processing. For this purpose, in the extended security mode, VC tables maintained by IP switches must be tagged to indicate if they belong to secure flows so that IP switches can send appropriate control protocol message to their upstream neighbors if a VC secure flow is reclaimed or lost due to switch or link failure in a downstream neighbor.

5 Conclusion

Both high speed and security in networking are expensive features to implement and organizations willingness to bear this cost is tied to the new applications enabled by the availability

of secure high speed networks. Corporations that are transforming their business processes to take full advantage of on–line electronic commerce capabilities are the most likely candidates to invest in this technology. It is with this view that we have focused our attention on providing integrated high speed and security for extranets. Security and high speed are properties that are generally orthogonal to each other. This is mainly due to the heavy computational costs associated with the cryptographic primitives use to implement security mechanisms. Therefore, in secure high speed networking we aim to increase network speed while reducing the overall cryptographic computational costs.

The functional integration of firewall and IP switch outlined in section 4 require further investigation to address many operational issues that may arise due to the modification required in existing protocols. A significant disadvantage of designing a security network system optimized for speed by utilizing non–standard techniques is the resultant loss of generality for the entire system. However, this would be a carefully considered implementation decision a system administrator will have to make when building a high speed extranet.

References

[1] R. Atkinson. *IP Authentication Header (AH)*. IETF RFC 1826, Aug 1995.

[2] R. Atkinson. *IP Encapsulating Security Payload (ESP)*. IETF RFC 1827, Aug 1995.

[3] R. Atkinson. *Security Architecture for the Internet Protocol*. IETF RFC 1825, Aug 1995.

[4] H. Esaki, K. I. Nagami, and M. Ohta. High speed datagram delivery over internet using ATM technology. In *Proceeding of the Networld+Interop*, Las Vegas, NV, Mar 1995.

[5] R. Gareiss. Industrial–strength extranet. *Data Communications Magazine*, pages 71–82, Jun 1997.

[6] S. Lin and N. McKeown. A simulation study of IP switching. In *Proceeding of the ACM Sigcomm*, Cannes, France, Sep 1997.

[7] P. Newman, W. L. Edwards, R. Hinden, E. Hoffman, F. C. Liaw, T. Lyon, and G. Minshall. *Ipsilon's General Switch Management Protocol Specification Version 1.1*. IETF RFC 1987, Aug 1996.

[8] P. Newman, W. L. Edwards, R. Hinden, E. Hoffman, F. C. Liaw, T. Lyon, and G. Minshall. *Ipsilon Flow Management Protocol Specification for IPv4*. IETF RFC 1953, May 1996.

[9] P. Newman, T. Lyon, and G. Minshall. Flow labeled IP: A connectionless approach to ATM. In *Proceedings of the IEEE Infocom*, pages 1251–1260, San Francisco, CA, Mar 1996.

[10] P. Newman, G. Minshall, T. Lyon, and L. Huston. IP switching and gigabit routers. *IEEE Communications Magazine*, pages 64–69, Jan 1997.

[11] R. Oppliger. Internet security: Firewalls and beyond. *Communications of the ACM*, 40(5):92–102, May 1997.

[12] G. Parulkar, D. C. Schmidt, and J. S. Turner. IP/ATM: A strategy for integrating IP with ATM. In *Proceedings of the ACM Sigcomm Symposium on Communications Architectures and Protocols*, Cambridge, MA, Sep 1995.

[13] J. Postel. *Internet Protocol*. IETF RFC 791, Sep 1981.

[14] Y. Rekhter, B. Davie, D. Katz, E. Rosen, and G. Swallow. *Tag Switching Architecture Overview*. IETF RFC 2105, Feb 1997.

[15] L. Zhang, S. E. Deering, D. Estrin, S. Shenker, and D. Zappala. RSVP: A new resource ReSerVation Protocol. *IEEE Network Magazine*, 9(5), 1993.

[16] Y. Zheng. Digital Signcryption or How to Achieve Cost(Signature & Encryption) $\ll$ Cost(Signature) + Cost(Encryption). In *Advances in Cryptology – Crypto'97*, Lecture Notes in Computer Science. Springer–Verlag, 1997.