

Design and Evaluation of a High-Performance ATM Firewall Switch and Its Applications

Jun Xu, *Student Member, IEEE*, and Mukesh Singhal, *Senior Member, IEEE*

Abstract—We present the design of a value-added ATM switch that is capable of performing packet-level (IP) filtering at the maximum throughput of 2.88 Gbit/s per port. This firewall switch nicely integrates the IP level security mechanisms into the hardware components of an ATM switch so that most of the filtering operations are performed in parallel with the normal cell processing, and most of its cost is absorbed into the base cost of the switch. The firewall switch employs the concept of “last cell hostage” (LCH) to avoid or reduce the latency caused by filtering. We analyze in detail the performance of the firewall switch in terms of the throughput and the latency and address related design issues. Applications of our firewall switch as Internet and intranet security solutions are also discussed.

Index Terms—ATM switch design, computer network security, firewalls, internetworking, multiprotocol over ATM (MPOA).

I. INTRODUCTION

A SECURITY hole of current asynchronous transfer mode (ATM) deployments is that an end-to-end ATM connection generally bypasses Internet and intranet firewalls, if there are any. Terminating ATM connections at an intermediate packet-filtering firewall for inspection is not a solution, as the throughput of the traditional router-based firewalls is still much smaller, compared to the ATM rate of OC-12c (622 Mbit/s) or higher. The overhead caused by segmentation and reassembly (SAR) may further bring the performance down. To circumvent this problem, the ATM Forum favors the avoidance of packet filtering by exerting discretions at the connection-establishment time. Based on this policy, two alternative access-control schemes have been proposed. Smith [15] proposes that access-control decisions be made at connection-establishment time, based on the higher layer information (e.g., source and destination IP addresses, ports, etc.) that is contained in the ATM *signaling message* as *information elements*. After the connection is established, the access-control device “gets out of the way.” Obviously, this is unacceptable because an intruder can always lie about the service he wants to access at the connection-establishment time, and there is no way to check the contents of packets on a connection once the connection is established. Pierson [14] tries to fix this problem by proposing that “a new SVC (switched virtual connection) is requested when each new service is started.” However, this requires that the ATM layer be notified whenever a new socket is opened, which entails

considerable change to the whole TCP/IP stack and existing applications [5]. Even worse, this requires a new SVC for each transport layer flow, which may lead to the problem of *VC explosion*. Apparently *call screening* alone will not solve the problem of network-level access control in an ATM network. Therefore, a new ATM firewall architecture is called for that is capable of performing packet-level filtering at high speed.

At the time of writing, StorageTek’s ATLAS product was the only ATM packet-filtering device available in the literature [9]. ATLAS is a line filter that scans an ATM physical link to perform the packet-level filtering at the rate of OC-3 (150 Mbit/s). Two performance-boosting strategies are used in ATLAS. First, to avoid SAR, for each packet ATLAS only checks the first cell, which contains the IP header, TCP/UDP ports, protocol, and TCP flags (if applicable), to determine whether or not the packet is “safe.” If the packet is considered safe, all the following cells that belong to it are passed or otherwise dropped. Second, ATLAS uses a *policy cache architecture* [9] to dramatically speed up the process of deciding whether or not a packet header is safe. The core unit of this architecture is a cache block called the *policy cache*. Each entry of the policy cache is a combination of virtual path identifier/virtual connection identifier (VPI/VCI), source and destination IP addresses, and source and destination TCP/UDP ports that are considered “safe.” When the first cell of a packet arrives, it is compared with each entry in the policy cache. If a cache hit occurs, cells of the packet are forwarded. Otherwise, the first cell goes through a software-screening process, and all other cells of the packet are buffered in a queue. If the cell is found unsafe, the whole packet is dropped. Otherwise, the packet is allowed to pass, and an entry that contains the header pattern of the packet is added to the policy cache. The policy cache is implemented using content-addressable memory (CAM), which enables simultaneous searching of all memory locations to find a match with the pattern being searched for. Therefore, ATLAS can decide whether a packet header hits the policy cache very quickly.

However, the ATLAS product does have its limitations and drawbacks. First, it does not accept IP packets with IP option fields because an IP option can be as large as 40 bytes and may “push” the TCP header to the second cell. Second, it is not friendly to those who manage and administer it because TCP/IP filtering rules for any newly established PVC or SVC must be manually configured. This is impractical in future ATM networks where a large number of SVC’s will be established on-the-fly in a short period of time.

Manuscript received July 14, 1998; revised February 11, 1999.

The authors are with the Department of Computer and Information Science, Ohio State University, Columbus, OH 43210 USA (e-mail: jun@cis.ohio-state.edu; singhal@cis.ohio-state.edu).

Publisher Item Identifier S 0733-8716(99)04493-5.

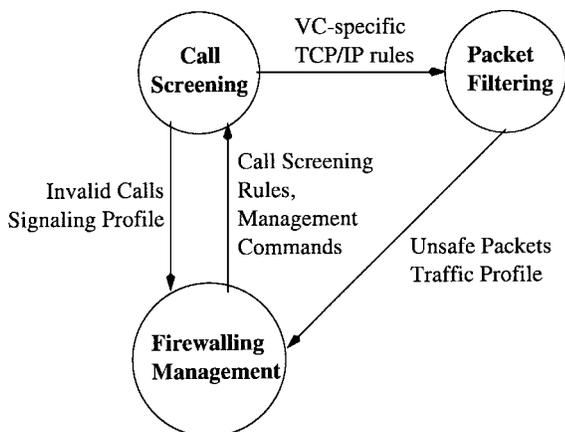


Fig. 1. Logical design of our firewall switch.

In this paper, we present an implementation-ready hardware design of a value-added ATM switch that performs packet-level filtering at high speed (2.88 Gbit/s) and present its applications as Internet and intranet security solutions. Sections II and III present, respectively, the logical and physical design of our firewall switch. In Section IV, we analyze the throughput and the latency of our firewall and explore related design issues. In Section V, we present applications of our firewall switch. Section VI concludes the paper.

II. LOGICAL DESIGN OF OUR FIREWALL SWITCH

As shown in Fig. 1, three related security services, namely, *call screening service* (CSS), *packet-filtering service* (PFS), and *firewall management service* (FMS), are implemented in our firewall switch. Their functions are described in detail in the following sections.

A. CSS

When an ATM *signaling message* that requests a new SVC to be established arrives at the firewall, CSS decides whether it is secure to establish the connection and/or whether the new connection needs to be filtered based on the identity and authentication information¹ contained in the signaling message. If the connection is considered secure, the firewall switch will allow its traffic to bypass packet-filtering mechanisms. We will see in Section III-A that implementation of such bypasses are trivial. For simplicity, in the following discussion we assume all connections that travel through the firewall switch need to be filtered. If the connection needs to be filtered, CSS will advise the PFS to filter the traffic in the connection and supply PFS with packet-filtering rules as shown in Fig. 1. In a firewall switch, filtering rules for various

¹Authentication of the signaling message is described in detail in [16].

connections may be different and may be generated on-the-fly, as follows. First, the firewall switch may determine source and destination IP addresses/prefixes, either from the higher layer information elements (IE's) contained in the authenticated signaling message or from the source and destination ATM addresses, by querying a preconfigured ATM address to the IP address/prefix mapping table. Then, from the intended services indicated in the signaling message or in the security policy, the firewall switch may determine the source and the destination TCP/UDP ports allowed in the connection. Finally, these two pieces of information are glued together to generate a set of filtering rules. For example, if the source is a router for subnet 164.107.*.*, and the destination is a router for subnet 192.41.245.*, and the destination subnet allows only HTTP (port 80) and TELNET (port 23) from the source subnet, then the TCP/IP filtering rules are as shown at the bottom of the page.

Filtering rules for securing reverse traffic will be generated in the same fashion. A default set of TCP/IP filtering rules will be used by a PFS if CSS does not have enough information to generate such filtering rules on-the-fly. Such automatic configuration of TCP/IP filtering rules is vital in future ATM environments where a large number of SVC connections will be established within a short period of time. This relieves the network managers from worrying about setting up filtering rules for each SVC when it is established.

B. PFS

PFS inspects the header of IP packets to block “unsafe” packets while allowing the “safe” packets to pass. We use a scheme called last cell hostage (LCH) [18] in our firewall switch to reduce the latency incurred by packet filtering. Two performance-boosting schemes used in ATLAS, namely, filtering only on the first cell and hardware caching of the security policy (see Section I-B), are also used in our firewall switch.

In ATLAS, if the first cell misses the policy cache, all cells of the packet have to be buffered while waiting for a slow software inspection process to finish (even when the last cell of the packet arrives later than the inspection decision). In contrast, with our LCH scheme, all cells of a packet except the last one are allowed to pass, even if a cache miss occurs. Only the last cell is kept as a “hostage” before the inspection is finished. If the inspection finds out that the packet is safe, the last cell is passed. Otherwise, it is dropped and is substituted with a pseudo-last cell whose payload is generated randomly so that the cyclic redundancy code (CRC) failure of the whole packet at the receiver is guaranteed. This prevents the drop of one packet from affecting the following packet, which will otherwise be mixed with the previous packet and cause

Src IP	Dest IP	Src Port	Dst Port	Protocol	Action
164.107.*.*	192.41.245.*	>1023	80	TCP	Allow
164.107.*.*	192.41.245.*	>1023	23	TCP	Allow
(All other packets are blocked)					

corruption. In AAL5, the last cell of a packet is explicitly identified by a 3-bit payload type identifier (PTI) field in the header, and the receiver will not reassemble a packet from the cells without this last cell. Therefore, this CRC-failure enforcement cannot be bypassed by a malicious packet. The use of LCH in our firewall switch has the following two advantages.

- Even when a packet header misses the policy cache, if the packet is reasonably long, the software-based filtering process can be finished before the last cell arrives. This means that no latency is experienced even when a cache miss occurs. Section IV-B2 will demonstrate LCH's effectiveness in avoiding or reducing latency caused by packet filtering.
- LCH makes it easy to process IP packets with options. The technical challenge in filtering IP packets with an option is that the decision may not be made until the second cell arrives. ATLAS does not process such packets because it would have to consider each of such a packet as missing the policy cache. When the percentage of such packets is high, the processing overhead would be unbearable in ATLAS. However, with LCH, only some fields of the first cell need to be kept as the state information, and the software-based filtering process is postponed until the second cell arrives. Since all the cells except the last one can be passed without delay, the IP packets with options can be processed almost as fast as those without.

Even though the LCH scheme only blocks the last cell of a packet while unconditionally allowing other cells to pass, it is by all means a safe approach. If an attacker sends a "bad" packet to the system, the packet will be discarded by the receiver because its last cell will be "corrupted."

C. FMS

FMS controls and manages CSS and PFS and provides user-friendly administration tools to network managers. These tools allow network managers to perform operations such as specifying or updating call screening rules, monitoring abnormal user activities, and disabling connections that violate the security policy. FMS logs two types of events forwarded from CSS and PFS. The first type of event includes unsafe signaling messages detected by CSS and unsafe packets detected by PFS. Such information can be used to analyze what type of intrusion is involved and how security policy needs to be updated in response to the intrusion. The other type of event includes the signaling and traffic profile information about each connection. Such information allows our firewall switch to detect abnormal user activities that may indicate an intrusion attempt.

III. THE PHYSICAL DESIGN OF OUR FIREWALL SWITCH

Our firewall switch modifies various components of a normal ATM switch to implement the security services described in Section II. Conceptually, a normal ATM switch consists of five functional modules, namely, input module (IM), output module (OM), cell switching fabric (CSF), call admission con-

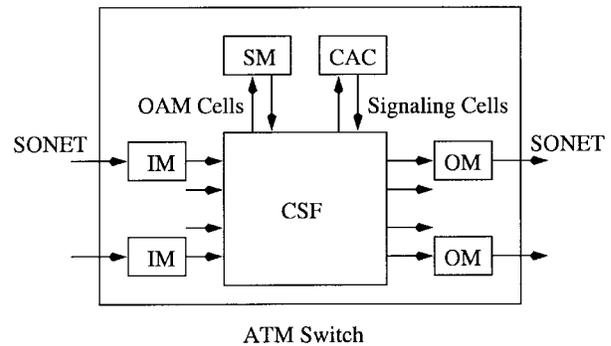


Fig. 2. The conceptual organization of an ATM switch [3].

rol (CAC), and system management (SM), as shown in Fig. 2. IM extracts cells from the incoming SONET payload envelope, determines the destination port of each cell, and writes it to an internal tag appended to the cell. CSF routes each cell to its destination port which is indicated in the internal tag. CAC assembles the signaling cells into a signaling message and processes it. SM processes the operations, administration, and maintenance (OAM) cells and performs switch-management functions. All outgoing cells arrive at OM through CSF. OM processes and removes the internal tags appended to the cells and puts the cells into the SONET payload envelopes for transmission. Our firewall switch nicely integrates the IP-level security mechanisms into these hardware components of an ATM switch so that most of the filtering operations are performed in parallel with the normal cell processing, and most of its cost is absorbed into the base cost of the switch. In our firewall switch, IM and OM are enhanced to implement PFS, CAC is enhanced to implement CSS, and SM is enhanced to implement firewall-management service. In Sections III-A–III-D, we discuss the modifications to IM, OM, CAC, and SM, respectively.

A. IM

In IM, translation of the VPI/VCI of each incoming cell to determine its output VPI/VCI/port is performed in a functional block called *cell translation*. In a firewall switch, this block is enhanced to perform filtering operations as well. Fig. 3 shows the overall picture of the cell-translation block in a firewall switch. Area A depicts the header-translation process in a normal switch. In flow 1, VPI/VCI of an incoming cell (stored in the cell buffer) is used as the key to search in the VP/VC table. The corresponding output VPI/VCI/port is fed to normal header-translation logic (flow 2), which will in turn update the VPI/VCI field of the cell buffer (flow 3) and write the output-port information to the internal tag (flow 4).

In Fig. 3, Area B shows the packet-filtering function in a firewall switch. The VP/VC table is enhanced to contain the following firewall-related fields. EXP-1st and EXP-2nd indicate whether the incoming cell is the first or second cell of the packet. PASS and LCH indicate whether the packet should be passed or put into LCH, respectively. SNO is the serial number assigned to each packet. SNO_LCH and output module check (OMCHK) will be explained in Section III-B. If the firewall switch needs to distinguish whether a connection

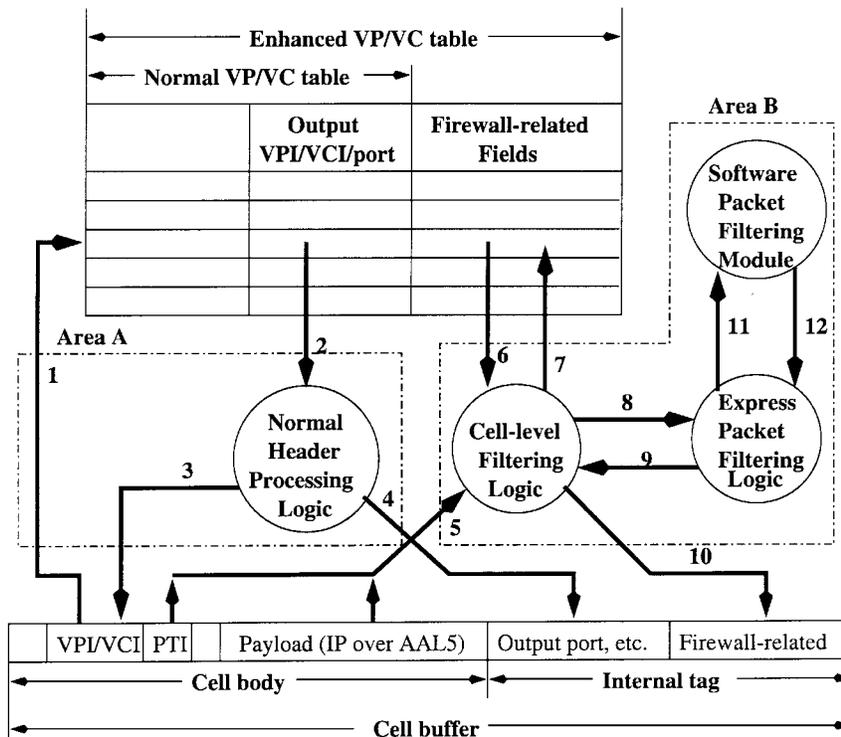


Fig. 3. The header translation module in our firewall switch.

needs to be filtered, a one-bit field can be added to VP/VC table for this purpose. These fields keep track of cell-level and packet-level state information regarding the connection. These fields (flow 6), PTI, and fields in IP payload (flow 5) will be fed to the *cell-level filtering logic* (CFL). CFL will decide whether the packet should be passed or put into LCH. It then updates the firewall-related fields in VP/VC (flow 7) and writes information to the firewall-related fields in the internal tag (flow 10 in Fig. 3).

Fig. 4 shows the flowchart of the CFL. If the cell is a nonfirst cell of a packet, then decision on the packet has already been made and is recorded in the PASS (to pass the packet) and LCH (to put the cell into LCH), respectively. Otherwise (“EXP-1st?”), CFL sends this first cell to express packet-filtering logic (EPFL) for further inspection (flow 8 in Fig. 3) and waits for a reply (flow 9 in Fig. 3). The decision is either to pass the packet or to put the packet into LCH, which is recorded in the PASS or LCH field (set to TRUE), respectively. Besides, if the decision is to put the packet into LCH, the OMCHK bit should be turned to TRUE and SNO.LCH will record the serial number of this packet. We will explain their use in Section III-B. If the packet contains an IP option, EXP-2nd is set to TRUE to indicate that the second cell is expected for inspection. When such a cell arrives (“EXP-2nd?”), it will be forwarded to EPFL as well. When the last cell of a packet arrives (“Last cell?”), as indicated by the PTI field in the cell header, the DROP, LCH, EXP-1st, and EXP-2nd will be reset or updated to “welcome” the arrival of the next packet. The value of LCH, OMCHK, and SNO will be carried to OM by T-LCH, T-OMCHK, and T-SNO fields in the internal tag, respectively.

EPFL employs a policy cache to boost the filtering performance. The policy cache stores the $\langle \text{VPI/VCI, Src IP, Dst IP, source port, destination port, and protocol} \rangle$ tuples that are found secure. When a cell is forwarded from CFL, it is checked against the cache block for a match. If a cache hit occurs, the decision “PASS” is sent to CFL. If a cache miss occurs, the decision “LCH” is sent to CFL, and the cell is forwarded to the software packet-filtering module (SPFM) for further inspection (flow 11 in Fig. 3). SPFM checks the cell (in software) against packet-filtering rules to see if it is secure. Once the final decision regarding whether the packet should be passed or dropped is available from SPFM (flow 12 in Fig. 3), EPFL will forward the decision to OM using internal cells or other internal communication mechanisms. Besides, a tuple that consists of the corresponding header fields of the packet will be added to the policy cache. For packets that contain an IP option, the first cell may be stored as the state information if needed.

The throughput of CFL/EPFL can be estimated as follows. We assume that the VP/VC table is implemented using 50 ns off-the-shelf DRAM. With off-the-shelf CMOS technology, CFL can be implemented within one memory cycle because its logic is very simple. We will show in Section IV-B1 that the length of a typical IP packet today is eight cells. We conservatively estimate the upper bound of the time that needs to be spent on such an average-length packet as follows. A read from the VP/VC table is needed before processing each and every cell, and a write to the VP/VC table is needed after processing the first, the second (when the packet contains an IP option), and the last cell. Processing the first cell also involves a lookup in a content-addressable memory module (the policy

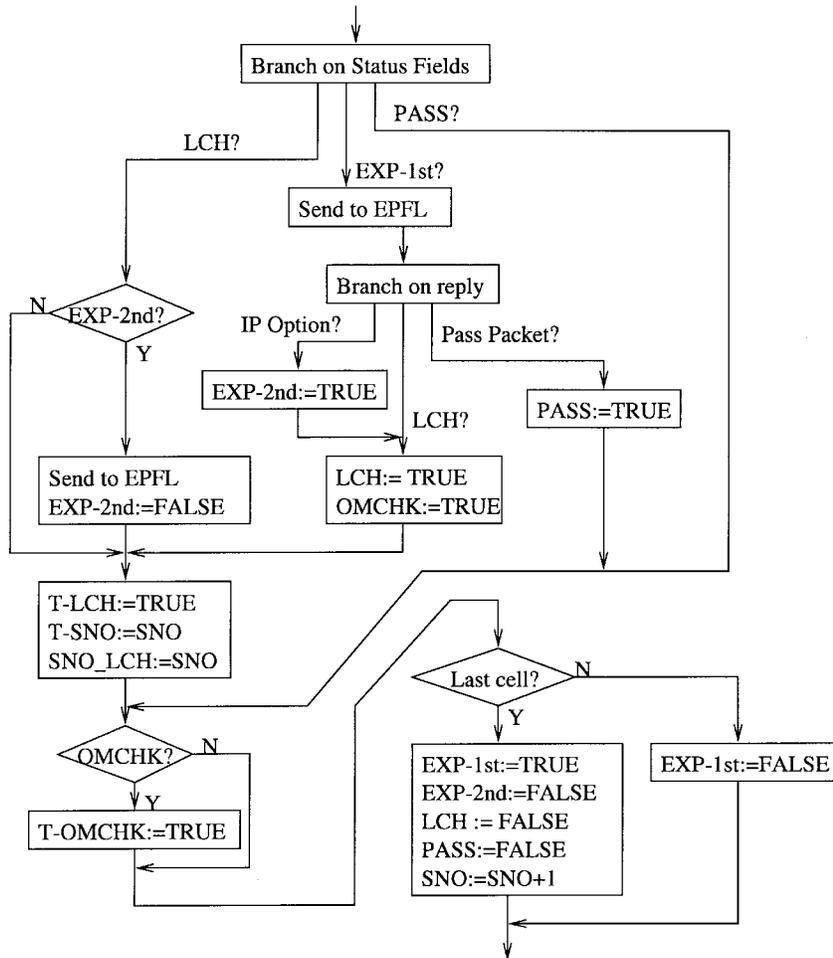


Fig. 4. The flowchart of CFL.

cache) in EPFL, which is conservatively estimated to be 150 ns.² Even without any pipelining, the first, second (possibly), and the last cell take, respectively, a maximum of six, three, and three memory cycles to process. Every other cell takes two memory cycles. Therefore, the average time spent on each cell is $(6+3+2+2+2+2+2+3)/8 = 2.75$ (memory cycles). At this rate, EPFL can process at least 7.27 million cells/s, which is equivalent to 3.08 Gbit/s. We will see in Section IV-B1 that this throughput is greater than the estimated throughput of the firewall switch, which is 2.88 Gbit/s. Thus, CFL/EPFL is not the bottleneck in our firewall switch.

B. OM

In a standard switch, OM processes the internal tag of each cell for housekeeping purposes and updates output VPI/VCI for multicast cells. OM also contains a VP/VC table, which contains fields for housekeeping information, such as the number of cells transported in each connection and a field that stores output VPI/VCI for multicast connections.

In a firewall switch, OM is involved in implementing the LCH scheme. When the LCH scheme needs to be applied to a

packet, IM will indicate that in the T-LCH field of the internal tag. Decision on that packet will also be forwarded to OM as soon as it is available. The responsibility of OM is to keep the last cell of the packet as a “hostage” until the decision on the packet arrives. To accomplish this function, the VP/VC table is enhanced to include two new fields. One is LCH_Cell, which points to the buffer that stores the last cell of an LCH packet and other cells that are blocked due to the LCH cell. The other is SEL_NO, which stores the serial number of the last decision received (PASS or DROP). Also, there is a decision queue that stores the decisions forwarded from IM’s, yet to be applied to LCH packets at OM.

The following processes are involved in implementing the LCH scheme. The first process is invoked when the last cell of an LCH packet arrives. The second process is invoked when other packets that follow the LCH packet in the same connection arrive. The third process is invoked when a decision on an LCH packet arrives.

- 1) OM identifies the last cell of an LCH packet from the LCH bit in the internal tag and the PTI field in the cell header. When such a cell is identified, OM will check whether the decision on the packet has arrived ($SEL_NO \geq T_SEL_NO?$). If the decision is available, it is retrieved from the decision queue and is applied to the cell.

²Normal access time for current 1k-size CAM is between 50–90 ns. However, larger-size CAM (e.g., 64 k) that may be needed by our firewall switch will have long access time due to propagation delay of OR gates used to concatenate small-size CAM chips into a large-size CAM.

- 2) While the decision on the LCH packet is pending, packets that follow the LCH packet in the same connection should also be buffered in order to preserve the cell order. This is taken care of by the T-OMCHK bit in the internal tag of a cell. IM will turn the T-OMCHK bit on once an LCH decision has been made to a packet in that connection. When OM detects such a cell, OM will check whether there is a cell buffered in LCH_Cell. If there is none, the cells should be allowed to pass. Otherwise the cells need to be queued up.
- 3) When the decision on a packet arrives, OM puts it in the decision queue and updates the SEL_NO field with the serial number contained in the decision message. Also, cells queued up in that connection will now be processed by OM.

The overhead of the first and third process can be absorbed into the standard housekeeping process of OM. The overhead incurred by the second process is equivalent to the processing of multicast cells because both of them involve a VPI/VCI table lookup.

As each OMCHK cell results in as much overhead as a multicast cell, it is desirable to make the number of such cells as small as possible. In the ideal situation, once a queue for OMCHK cells in a VP/VC becomes empty, the OMCHK bit is cleared in the corresponding VP/VC entry in IM. This is accomplished by OMCHK clearance protocol, which works as follows: once a queue becomes empty, an internal cell is sent back to IM, which contains the VPI/VCI of the connection and the serial number of the last packet that is put into LCH and has been processed by OM (recorded in SEL_NO). When IM receives the internal cell, it checks whether the serial number matches with the SEL_LCH field in the enhanced VP/VC table. If it matches, this means that no new LCH packet has been sent to OM after that one, and IM can safely turn the OMCHK bit off. Otherwise, this message is ignored.

C. CAC

CAC in a standard switch processes and routes signaling messages. In a firewall switch, CAC implements the CSS, as discussed in Section II-A. It compares the identities of the communicating parties contained in the signaling message with call screening rules to decide whether the connection is allowed to be established and/or whether the connection needs to be filtered. If the connection needs to be filtered, CAC will set up an entry in the VP/VC table for the new connection and initialize firewall-related fields in the VP/VC table. It will also generate the packet-filtering rules to be executed by SPFM. CAC may be assisted by fast cryptographic hardware [6] to quickly authenticate the digital signature contained in the signaling message, if applicable.

D. SM

In a standard switch, SM handles all the management-plane functions, including fault, performance, configuration, accounting, security, and traffic management [3]. In a firewall switch, it manages firewall functions that need to be performed by the switch, which involves the following tasks:

- maintaining firewall-related *managed objects*, such as call screening rules;
- executing the commands sent from the network management station, such as updating call screening rules;
- monitoring the performance of packet filtering, such as the throughput and average latency of filtering operation at each port.

IV. THE PERFORMANCE OF THE FIREWALL SWITCH

In this section, we study two important performance metrics of our firewall switch: throughput and latency. Related design issues such as the size of the policy cache and its replacement policy will also be addressed. In Section IV-A, we first present the statistical model and mathematical analysis for calculating the throughput and the latency of the firewall switch. In Section IV-B, data from real-world measurements and simulation are plugged into the expressions to obtain the numerical results.

A. Performance Analysis

1) *Throughput of the Firewall Switch*: Given a certain traffic load, the throughput of the firewall switch THR can be expressed as

$$\text{THR} = \frac{\text{STHR} * \text{NC}}{\text{MR}} \quad (1)$$

where MR is the miss ratio of the policy cache under the traffic load, NC is the average number of cells that a packet consists of in the traffic load, and STHR is the throughput of SPFM in IM. NC can be directly measured from the traffic sample and STHR can be estimated using the throughput data of traditional IP firewalls. MR, however, merits more discussion.

MR is closely related to the concept of *flow*. A *flow* is a unidirectional stream of IP packets with common tuple {source-IP-address, source-port, destination-IP-address, destination-port, protocol} in the header. A flow is started when the first packet of the flow arrives and is considered *expired*, i.e., when there has been no activity for a timeout period D_{expire} (typically set to 64 s [4], [17]). This concept of flow is actually an extension of Jain's *packet train model* [7] at layer-4, and D_{expire} is a counterpart of maximum allowed inter-car gap (MAIG) in that model. A flow is said to be *active* from the time it was created until the time it expires. When the first packet of a flow arrives at the firewall switch, it will obviously miss the policy cache and will be inspected in SPFM. If it is found secure, a cache entry that contains the tuple will be added to the policy cache. If the cache entry is not evicted from the policy cache before the last packet of the flow arrives, there is exactly one cache miss for all the packets inside the flow. If this is true for each and every flow, then $1/\text{PF}$ will be the miss ratio of the cache, where PF is the average number of packets per flow and can be directly measured from the traffic load. In the following sections, we show that a miss ratio of only 10% higher than $1/\text{PF}$ can be achieved with a "decent amount of cache" and the first-in/first-out (FIFO) cache replacement policy.

2) *On the Size of Policy Cache:* In this section, we explain what a “decent amount of cache” is and what happens if the cache size is less than that. Given a certain traffic load, we use $F(t)$ to denote the number of *active flows* at time t during an observation period $[0, T]$. Let $\text{MAXF} = \max_{0 \leq t \leq T} F(t)$ be the maximum number of active flows during the observation period $[0, T]$. We observe from National Laboratory for Applied Network Research’s (NLNR’s) statistics [4], [17] that the number of active flows $F(t)$ is quite stationary over time, given that the traffic volume is quite stationary. Therefore, MAXF should be very close to the average number of active flows. Let CS denote the size of policy cache in our firewall switch. If $\text{CS} \geq \text{MAXF}$, with the least recently used (LRU), the miss ratio of no more than $1/\text{PF}$ can be achieved. The reason is that whenever the first packet of a flow arrives, since the number of active flows at that moment is no more than MAXF , there must be one flow that is no longer active and whose header is present in the policy cache. Clearly, this entry can be evicted without causing a cache miss. The victim entry in an LRU cache clearly is such an entry. However, LRU is extremely expensive to implement in hardware. Realistic cache replacement policies like FIFO and RAND will have to be used.

On the other hand, if $\text{CS} < \text{MAXF}$, we say that a *cache overflow* happens at time t when $\text{CS} < F(t)$. At such a time, entries that cache the headers of active flows have to be evicted. We define *overflow ratio* (OR) as the average percentage of cache overflows relative to the number of cache entries. From a trace-driven simulation, we found that a large percentage of OR will contribute to the total miss ratio. Therefore, we define a “decent amount of cache” as a size larger than MAXF . Additionally, an even larger cache size helps reduce the cache miss ratio if the cache replacement policy used is not as good as LRU.

3) *Latency of the Firewall Switch:* In this section, we demonstrate the effectiveness of LCH in avoiding and reducing latency that is caused by packet-level filtering. Let us consider the following scenario. Assume that the total volume of the link is C (cells/s), and we would like to investigate how a packet on a connection of average bandwidth B (cells/s) will be delayed assuming that it misses the policy cache. The following three metrics are studied:

- $\text{Pr}(DE_{\text{withLCH}} = 0)$, the probability that no latency is experienced by the packet (such probability is always zero if LCH is not used because packet filtering in software takes a while);
- $\text{EXP}[DE_{\text{withLCH}}]$, the average latency experienced by the packet if LCH is used;
- $\text{EXP}[DE_{\text{w/oLCH}}]$, the average latency experienced by the packet if LCH is not used.

We assume that cells of different connections are fully interleaved. During each cell slot, the probability λ that an incoming cell belongs to the connection can be expressed as $\lambda = B/C$. Let T denote the interarrival time (in the unit of cell slots) between two consecutive cells in the connection, then $\text{Pr}[T = i] = \lambda(1 - \lambda)^{i-1}$, $i = 1, 2, 3, \dots$, and $\lambda = B/C$. When λ is small, this discrete geometric distribution can be

approximated as the exponential distribution $\text{Pr}(T > t) = \lambda e^{-\lambda t}$ [11]. Suppose the first cell of the packet arrives at time 0, the *response time* of the SPFM is R cell time, and N is the number of cells in the connection that arrive within time window $(0, R]$, then N is a random variable with Poisson distribution $\text{Pr}[N = k] = (\lambda R)^k e^{-\lambda R} / k!$. No latency is experienced by the packet that consists of L cells if no more than $L - 2$ cells arrive within time window $(0, R]$. Therefore

$$\begin{aligned} \text{Pr}[DE_{\text{withLCH}} = 0] &= \text{Pr}[N \leq L - 2] \\ &= \sum_{k=0}^{L-2} \frac{(\lambda R)^k e^{-\lambda R}}{k!}. \end{aligned} \quad (2)$$

We define the latency experienced by a packet as the maximum among the latency experienced by individual cells of the packet. With LCH, the cell that was delayed most is the last cell; without LCH, the cell that was delayed most is the first cell since it has to wait for the last cell to arrive. Therefore

$$\begin{aligned} DE_{\text{withLCH}} &= \begin{cases} R - \hat{T}: & R \geq \hat{T} \\ 0: & R < \hat{T} \end{cases} \\ DE_{\text{w/oLCH}} &= \begin{cases} R: & R \geq \hat{T} \\ \hat{T}: & R < \hat{T} \end{cases} \end{aligned} \quad (3)$$

where \hat{T} is the arrival time of the L th cell, assuming the arrival time of the first cell is zero. The density function of \hat{T} is

$$f_{\hat{T}}(t) = f_{ER_{L-1}, \lambda}(t) = \frac{\lambda^{L-1} t^{L-2} e^{-\lambda t}}{(L-2)!} \quad (4)$$

where $ER_{n, \lambda}$ is the Erlang random variable [11]. Therefore

$$\begin{aligned} \text{EXP}[DE_{\text{withLCH}}] &= \int_0^R (R - t) f_{\hat{T}}(t) dt \\ &= R \left(1 - \sum_{k=0}^{L-2} \frac{(\lambda R)^k e^{-\lambda R}}{k!} \right) \\ &\quad - \frac{L-1}{\lambda} \left(1 - \sum_{k=0}^{L-1} \frac{(\lambda R)^k e^{-\lambda R}}{k!} \right) \end{aligned} \quad (5)$$

$$\begin{aligned} \text{EXP}[DE_{\text{w/oLCH}}] &= \int_0^R R f_{\hat{T}}(t) dt + \int_R^\infty t f_{\hat{T}}(t) dt \\ &= R \left(1 - \sum_{k=0}^{L-2} \frac{(\lambda R)^k e^{-\lambda R}}{k!} \right) \\ &\quad + \frac{L-1}{\lambda} \left(\sum_{k=0}^{L-1} \frac{(\lambda R)^k e^{-\lambda R}}{k!} \right). \end{aligned} \quad (6)$$

B. The Numerical Results

In this section, we calculate numerical results of the firewall-switch performance by plugging in the parameters measured from the Internet backbone traffic [4]. The performance analysis of our firewall switch owes thanks to NLNR’s Internet traffic raw trace and statistics data (available online) [4]. Though their data was collected for studying unrelated issues such as TCP implementation behavior, routing dynamics, and

web caching, it is also quite useful in the quantitative estimate of the throughput and latency of our firewall switch. NLNR logged and collected statistics on traffic parameters transferred through the Internet and campus backbones, among which PF, the average number of packets in a flow, and PL, the average length of a packet in bytes, will be used in our analysis. Since only less than 1% of the traffic is not end-to-end (mostly traffic between backbone routers), it is very representative of the traffic that flows between the Internet and an enterprise network. Therefore, the traffic statistics are suitable for analyzing the firewall-switch performance.

1) *Throughput of the Firewall Switch:* We are now ready to calculate THR using (1). According to NLNR’s traffic statistics, $PL \approx 343$ (bytes) and $PF \approx 14.92$ (packets/flow). Then calculated from PL, $NC \approx 8$ (cells/packet) based on AAL 5 encapsulation. A survey on the performance of the state of the art software-based IP firewall technology [8], [12], [13] shows that $STHR \approx 50\,000$ (packets/s). What is still missing from (1) is MR. As we discussed in Section IV-A1, the value of MR depends on the size of the policy cache and the cache replacement policy being used. We conducted a simulation on the NLNR’s Internet trace data to examine the miss ratio of FIFO. We found that when CS is 20% larger than MAXF, FIFO can achieve a miss ratio that is only 10% larger than $1/PF$. Since implementing FIFO (at virtually no cost) involves much less cost than implementing LRU, it is chosen as the cache replacement policy. Therefore, $MR \approx 0.0737$ considering the 10% degradation due to the use of FIFO cache replacement policy. According to (1)

$$\begin{aligned} THR &= \frac{STHR * NC}{MR} \\ &= \frac{50\,000 * 8}{0.0737} \\ &= 5.43 * 10^6 \text{ (cells/s)} \\ &= 2.88 \text{ Gbit/s.} \end{aligned} \tag{7}$$

Therefore, our firewall switch can achieve a maximum throughput of 2.88 Gbit/s based on the current Internet traffic characteristics. As discussed in Section III-A, CFL/EPFL will not prevent this throughput from being achieved.

2) *Latency of the Firewall Switch:* In Section IV-A3, we derived the expressions to calculate three metrics. In this section, we plug in real-life data to get the numerical results. We first need to calculate the response time of SPFM R [used in (6) and (7)]. Let ATR be the actual throughput of the firewall switch. Since SPFM is the system bottleneck, the system utilization of SPFM is $\rho = ATR/STHR$. Response time R , which is the interval between the time a packet (the first cell) that misses the policy cache is sent to the SPFM and the time when the inspection finishes, can be expressed as (assuming an M/M/1 queue [11]) $R = PT/(1 - \rho)$, where $PT = 1/STHR$ is the processing time of a packet at SPFM.

There are two types of traffic which may miss the policy cache. The first type is the first packet of a TCP flow. In this case, it is either a 44-byte-long SYN packet (inverse direction) or a 40-byte-long dataless ACK (reverse direction), which consists of only two cells. The second type is the first packet

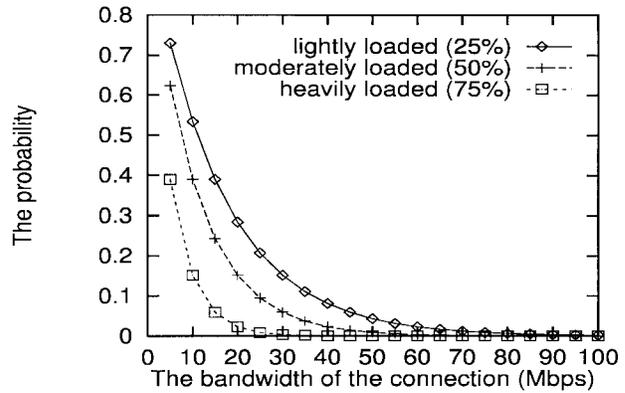


Fig. 5. The probability that no latency is experienced by a two-cell-long packet.

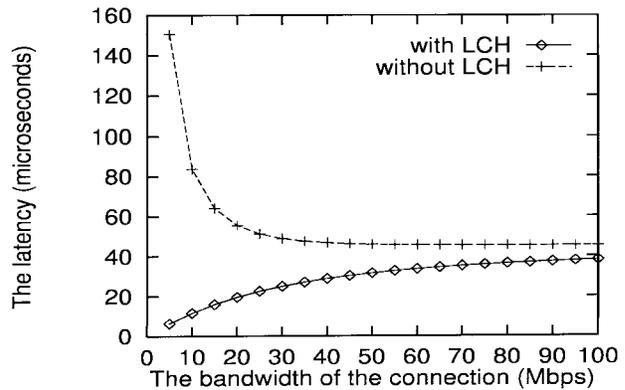


Fig. 6. The latency experienced by a two-cell-long packet under light (25%) load.

of a UDP flow or the nonfirst packet of a TCP flow that misses the policy cache due to cache replacement. The average length of such packets should be roughly equal to the average packet length of packets, which is eight cells long. In both cases, we discuss the latency caused by policy cache miss under heavy (75%), moderate (50%), and light (25%) load.

Fig. 5 shows $\Pr(DE_{withLCH} = 0)$ experienced by the first type of packet (two-cell-long) that misses the policy under the aforementioned three types of system loads. It shows that LCH is effective in avoiding latency when the system is lightly loaded or when the system is moderately loaded and the bandwidth of the connection is small (no more than the T3 rate). Figs. 6–8 compare $EXP[DE_{withLCH}]$ with $EXP[DE_{w/oLCH}]$ under heavy, moderate, and light loads, respectively. We can see that LCH is effective in avoiding or reducing latency when the system is lightly or moderately loaded or when the system is heavily loaded and the bandwidth of the connection is small. At the very least, LCH guarantees that the latency is less than R , which is around $80 \mu s$ under heavy load. As this type of packet is the first and second packet of a TCP connection, most probably the application on top of the connection was in the initialization stage (not in the middle of a high-resolution video conference session), and a small delay like this is generally tolerable.

Fig. 9 shows $\Pr(DE_{withLCH} = 0)$ experienced by the second type of packet (eight-cell-long) that misses the policy

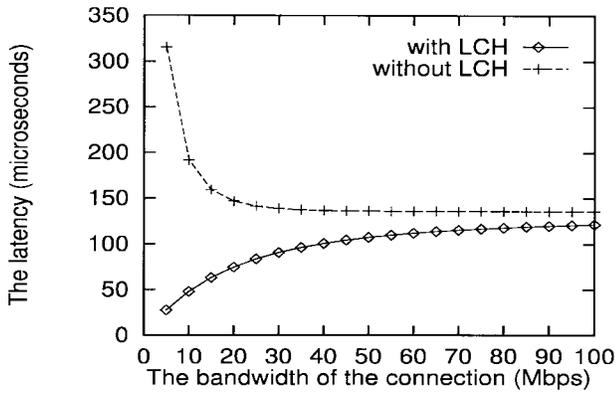


Fig. 7. The latency experienced by a two-cell-long packet under moderate (50%) load.

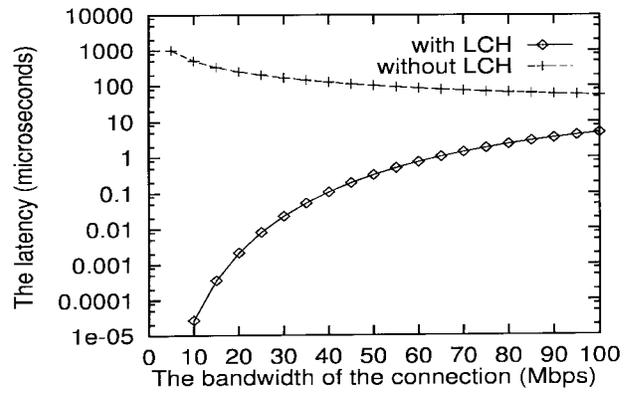


Fig. 10. The latency experienced an eight-cell-long packet under light (25%) load.

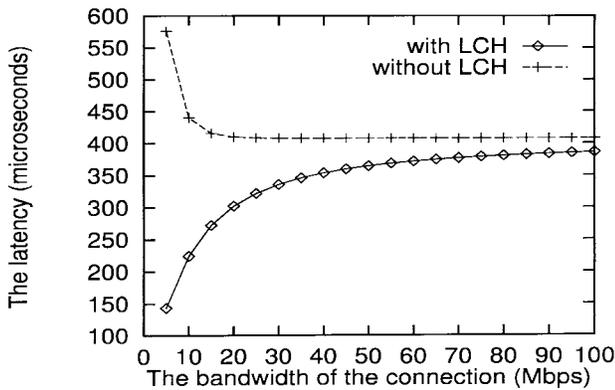


Fig. 8. The latency experienced by a two-cell-long packet under heavy (75%) load.

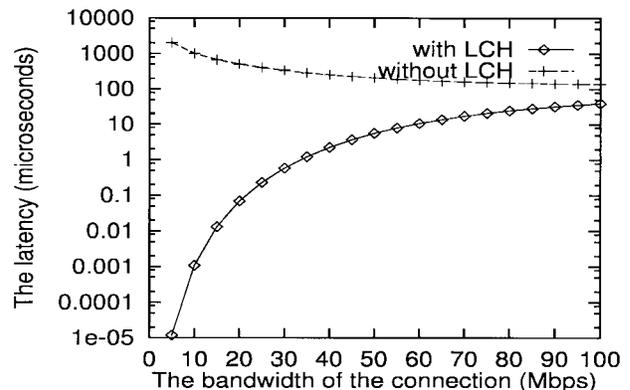


Fig. 11. The latency experienced by an eight-cell-long packet under moderate (50%) load.

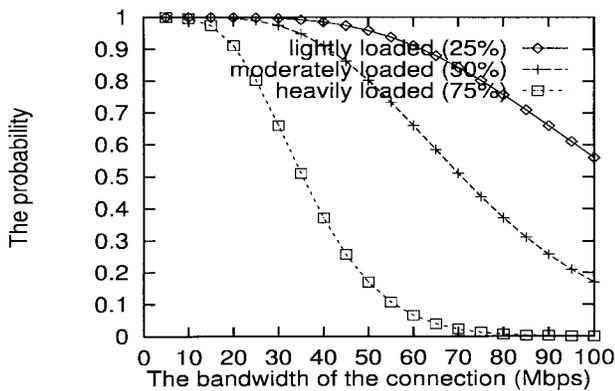


Fig. 9. The probability that no latency is experienced by an eight-cell-long packet.

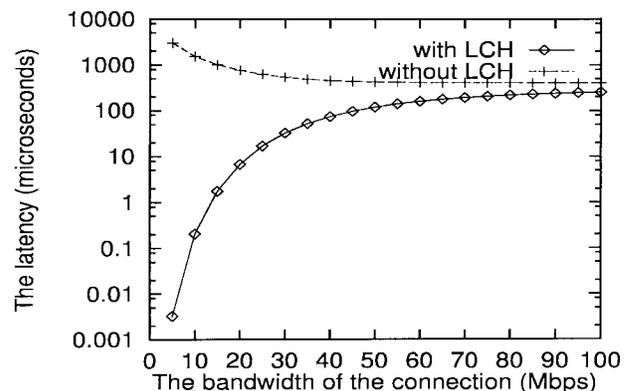


Fig. 12. The latency experienced by an eight-cell-long packet under heavy (75%) load.

under the three types of system loads. Figs. 10–12 compares $EXP[DE_{withLCH}]$ with $EXP[DE_{w/oLCH}]$ under heavy, moderate, and light loads, respectively (note that the Y axis is drawn on a logarithm scale). We see that LCH not only yields a high probability of avoiding latency, but also significantly reduces the average latency in all cases. Actually, it is this type of packet that needs such an improvement most, for two reasons. First, without LCH, packets as long as eight cells cause a significant latency when the bandwidth of the connection is small (e.g., more than 400 μs for a 10 Mbit/s

connection, as in Fig. 10) while LCH keeps the latency lower than R at all times. Second, such a packet may either carry a frame in a high-resolution video conference application (TCP) or a file block transferred using NFS (UDP). Delay of such packets is obviously undesirable, and LCH effectively avoids and reduces latency in such cases.

V. APPLICATIONS

An obvious use of the firewall switch is to place it at the “chokepoint” between an enterprise ATM network and a

public ATM WAN to secure all traffic that travels between the two networks. In the following, we show a novel use of our firewall switch to fix a security hole of multiprotocol over ATM (MPOA) [2] that has long been left unsolved. In MPOA, *cut-through connections* generally bypass firewall routers, if there are any, which may lead to security breaches [10]. The only solution proposed so far came from the ATM Forum. In this proposal, the first several packets between two end systems, which according to MPOA's flow detection process must travel through the *default route*,³ will get filtered by the intermediate MPOA servers. Only when these packets are found secure will the cut-through connection be allowed to be established. This approach is insecure because a malicious end system can always send "benign packets" in the beginning. Once the cut-through connection is established, it can do anything it pleases over that connection.

If there were a physical chokepoint inside the ATM network such that all cut-through connections must travel through it, then placing our firewall switch at that point would be a perfect firewall solution. However, generally there will not be such a physical chokepoint because a security domain [typically an IP subnet defined over a virtual LAN (VLAN)] can span multiple *edge devices* and an *edge device* can have ports that belong to different security domains. Instead, our firewall solution for MPOA is based on our novel concept of "logical chokepoint." In this approach, a *connection classification* protocol is executed between an edge device and an MPOA server to decide whether a cut-through connection is considered secure. If the connection is found insecure, it will be forced to travel through a firewall switch, which is deployed in the "middle" of an ATM network. In this way, participating edge devices act together to form a protection shield against internal security breaches. The *connection classification* protocol is nicely embedded into the MPOA address resolution protocol as follows. The MPOA *resolution request* that the source *edge device* sends to the MPOA server will contain the IP addresses of the source and the destination end systems and the authentication information of the source end system, such as a digital signature, if applicable. The MPOA server will decide whether the connection is secure based on this information. The decision will be included in the MPOA *resolution reply* sent from the MPOA server to the edge device. Forcing a connection to travel through a firewall switch can be implemented as follows. In the signaling message, the source-edge device specifies the ATM address of the firewall switch to be the final destination. The true destination is encoded into a user-to-user *information element* or *generic identifier transport information element* [1]. When this signaling message is forwarded to the firewall switch, the firewall switch knows the game being played, and it will retrieve the final destination from the corresponding information element and use it as the new destination to forward the signaling message. This scheme is presented in detail in [19]. We demonstrate in [19] that the only significant cost of the scheme is the firewall switch, which is amortized over the huge number of end systems it is able to protect. We

also show in [19] that the proposed scheme has no negative-performance impact on the edge devices and MPOA servers.

VI. SUMMARY

We designed a high-performance firewall switch that operates at the maximum throughput of 2.88 Gbit/s. Our firewall switch nicely integrates the IP-level security mechanisms into the hardware components of an ATM switch so that most of the filtering operations are performed in parallel with the normal cell processing, and most of its cost is absorbed into the base cost of an ATM switch. The proposed switch employs an LCH scheme to avoid or reduce latency caused by packet filtering. We conducted an in-depth performance analysis of the firewall switch in terms of throughput and latency and addressed related design issues. Finally, we proposed a novel use of our firewall switch to fix a security hole in MPOA that had been left unsolved.

ACKNOWLEDGMENT

The authors would like to thank anonymous referees for their constructive comments that helped to improve this paper.

REFERENCES

- [1] *Private network-network interface specification version 1.0 (PNNI 1.0)*. ATM Forum, Mar. 1996.
- [2] *Multi-protocol over ATM (MPOA) version 1.0*. ATM Forum, Feb. 1997.
- [3] T. Chen and S. Liu, *ATM Switching Systems*. Boston, MA: Artech House, 1995.
- [4] National Laboratory for Applied Network Research (NLNAR). (1998, June). *Flow statistics analysis for fix-west*. [Online]. Available www: <http://www.nlanr.net/>.
- [5] J. Hughes and A. Guha, "Requirements for secure packet-level access over ATM," ATM Forum/95-1126, 1995.
- [6] IBM Corp. (1998). *IBM 4758 PCI cryptographic coprocessor*. [Online]. Available www: <http://www.ibm.com/>.
- [7] R. Jain and S. Routhier, "Packet trains: Measurements and a new model for computer network traffic," *IEEE J. Select. Areas Commun.*, vol. 4, pp. 986-995, Sept. 1986.
- [8] Keylabs Inc. (1998, May). *Test final report—Firewall shootout network + interop*. [Online]. Available www: <http://www.keylabs.com/>.
- [9] B. Kowalski, "Atlas policy cache architecture." [Online]. StorageTek Corp. 1996. Available: <http://www.network.com/>.
- [10] D. Minoli and A. Alles, *LAN, ATM, and LAN Emulation Technologies*. Norwood, MA: Artech House, 1996.
- [11] R. Nelson, *Probability, Stochastic Processes, and Queueing Theory*. New York: Springer-Verlag, 1995.
- [12] D. Nessett and P. Humenn, "The multilayer firewall," in *Proc. NDSS '98*, San Diego, CA.
- [13] D. Newman, H. Holzbaur, and K. Bishop. (1997, March). Firewalls: Don't get burned. *Data Commun. Mag.* [Online]. Available: <http://www.datacomm.com/>.
- [14] L. Pierson and T. Tarman, "Requirements for security signalling," ATM Forum/95-0137, 1995.
- [15] T. Smith and J. Stidd, "Requirements and methodology for authenticated signalling," ATM Forum/94-1213, 1994.
- [16] T. Tarman, "Phase I ATM security specification," ATM Forum BTD-SECURITY-01.13, July 1997.
- [17] K. Thompson, G. Miller, and R. Wilder, "Wide-area internet traffic patterns and characteristics," *IEEE Network*, pp. 10-23, Nov. 1997.
- [18] J. Xu and M. Singhal, "Design of a high-performance ATM firewall," in *Proc. 5th ACM Conf. Comp. Commun. Security*, San Francisco, CA, Nov. 1998.
- [19] ———, "A firewalling scheme for securing MPOA-based corporate intranets," Department of CIS, Ohio State University, Columbus, OH. Rep. OSU-CISRC-6/98-TR19, 1998.

³MPOA trigger then needs to be used cautiously.



Jun Xu (S'98) received the B.S. degree in computer science from the Illinois Institute of Technology, Chicago, in 1995 and the M.S. degree in computer and information science from Ohio State University, Columbus, in 1997. He is currently a Ph.D. candidate in the Department of Computer and Information Science, Ohio State University.

His research interests include computer and communication security, protocol and architecture for high-speed networks, network performance modeling and simulation, and networking hardware design.



Mukesh Singhal (A'92–SM'98) received the B.S.E. degree in electronics and communication engineering (with high distinction) from the University of Roorkee, Roorkee, India, in 1980 and the Ph.D. degree in computer science from the University of Maryland, College Park, in May 1986.

He is an Associate Professor of Computer and Information Science at Ohio State University, Columbus. His current research interests include operating systems, distributed systems, mobile computing, high-speed networks, computer security, and performance modeling. He has published more than 100 refereed articles in these areas. He has coauthored two books, *Advanced Concepts in Operating Systems* (NY: McGraw-Hill, 1994) and *Readings in Distributed Computing Systems* (Piscataway, NJ: IEEE Computer Society Press, 1993). He is currently the Program Director of the Operating Systems and Compilers Program at the National Science Foundation.