

Network Firewall Policy Tries*

Errin W. Fulp and Stephen J. Tarsa
Department of Computer Science
Wake Forest University
Winston-Salem, NC, USA 27109
nsg.cs.wfu.edu
fulp@wfu.edu

Abstract

Network firewalls remain the forefront defense for most computer systems. These critical devices filter traffic by comparing arriving packets to a list of rules, or security policy, in a sequential manner. Unfortunately packet filtering in this fashion can result in significant traffic delays, which is problematic for applications that require strict Quality of Service (QoS) guarantees. Furthermore, as network speeds and capacities continue to increase, the processing time associated with filtering only compounds these delays. Given this demanding environment, new methods are needed to increase network firewall performance.

This paper introduces a new technique for representing a security policy in software that maintains policy integrity and provides more efficient processing. The policy is represented as an n -ary retrieval tree, also referred to as a trie. This structure is able to quickly reach decisions based on the security policy by simultaneously eliminating multiple rules with few comparisons. As a result, the worst case processing requirement for the policy trie is a fraction compared a list representation, which only consider rules individually (1/5 the processing for TCP/IP networks). Furthermore unlike other representations, the n -ary trie developed in this paper is proven to maintain policy integrity. The creation of policy trie structures is discussed in detail and their performance benefits are proven theoretically and validated empirically.

Keywords: Security management, firewalls, policy representation, and scalability.

1 Introduction

The benefits of highly interconnected computer networks have been accompanied by an increase in network-based security attacks. To address these threats, firewalls, also referred to as packet filters, have become a critical component of network security systems. Firewalls provide access control, auditing, and traffic control based on a security policy by inspecting packets sent between networks [3, 21]. The security policy is an ordered list of rules that defines an action to perform on arriving or departing packets. When a packet arrives at the firewall it is sequentially compared against the rules until a match is found [21]. This is referred to as a *first-match* policy and is used in the majority of firewall systems including the Linux firewall implementation `iptables` [17]. Once the match is determined the associated action is performed and the packet is either accepted or denied.

*This work was supported by the U.S. Department of Energy MICS (grant DE-FG02-03ER25581). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the DOE or the U.S. Government.

Firewalls are often implemented as a dedicated machine, similar to a router. Unfortunately, packet filtering requires a significantly higher amount of processing time than routing [16, 19]. Processing time increases as rule sets increase in length and complexity [4, 15]. As a result, a firewall in a high-speed environment (e.g. Gigabit Ethernet) can easily become a bottleneck and is susceptible to DoS attacks [4, 8, 11, 12]. These attacks merely inundate firewalls with traffic, delaying or preventing legitimate packets from being processed. Methods of improving firewall performance are needed for the next generation of high-speed networks and security threats.

Improving hardware is one method for increasing the amount of traffic a firewall can process. Current research is investigating different distributed firewall designs to reduce processing delay [4, 15] and possibly provide service differentiation [10]. While performance can increase using these methods, it requires multiple machines and/or specialized hardware. As a result these improvements are not amenable to legacy systems and thus do not provide a solution to many systems.

Improving software is another method to increase firewall performance that is applicable to a larger set of systems [5, 14, 16]. Similar to approaches that address the longest matching prefix problem for packet classification [6, 7, 9, 17, 18], these solutions employ better policy representations and searching algorithms. For example, retrieval trees (tries) offer quick search times and have been utilized to decrease the packet processing time [16, 18]. These policy models use the classical definition of a trie structure, which is a variation of a binary tree [1]. While this method groups rules in an efficient manner, the firewall tuples are stored in a binary format (one bit per branch) that increases the processing overhead and is difficult to implement (ultimately requiring a grid of binary tries) [7]. Furthermore, these binary trie structures are designed to determine the longest matching prefix, which results in the *best-match* rule (not typically used in network firewalls). As described in [18] first-match is possible, however it requires additional information and comparisons to rank possible rules, which is **not** required by the representation described in this paper.

In contrast, Directed Acyclical Graphs (DAG's) were used to store packet header fields (multibit field) in [6]. This structure was shown to efficiently store filter rules for layer four switching. However as described in section 3 of this paper, the DAG structure is unable to maintain integrity if partial-matching rules exist; thus, severely limiting its application to firewalls. Trees have also been successfully used to model firewall policies in [2, 13]; however, the primary purpose of this research was locating rule conflicts and anomalies, not improving processing time.

This paper introduces a new security policy representation called a policy trie, that is readily implemented and significantly reduces the packet processing time. Rules are represented as an ordered set of tuples, maintaining precedence relationships among rules and ensuring policy integrity (policy trie and list always arrive at the same result). The policy is modeled as an n -ary retrieval tree (trie), uniquely combining the retrieval efficiency of a trie and the flexibility of an n -ary tree.

When the policy trie is created rules are grouped by tuples (parts of the rule), allowing the elimination of multiple rules as a packet is processed and the trie is traversed. This is in contrast to the traditional list representation that can only consider one rule at a time. As a result, it will be proven that the policy trie has a worst case performance that is a fraction of a list representation ($1/k$, where k is the number of tuples). Furthermore unlike other representations, the policy trie maintains policy integrity; therefore this structure can easily and effectively represent current security policies. These theoretical results are verified via simulation under realistic conditions.

The remainder of this paper is organized as follows: Section 2 describes the models for firewall rules and a standard (list-based) security policy. The new policy representation called a policy trie

No.	Proto.	Source		Destination		Action
		IP	Port	IP	Port	
1	TCP	140.*	*	130.*	20	accept
2	TCP	140.*	*	*	80	accept
3	TCP	150.*	*	120.*	90	accept
4	UDP	150.*	*	*	3030	accept
5	*	*	*	*	*	deny

Table 1: Example TCP/IP security policy consisting of multiple ordered rules.

is introduced in section 3. Methods for creating and searching the policy trie, and proofs for policy integrity and performance bounds are also presented. Storage requirements for the representations are described in section 4. The performance of the policy trie is investigated experimentally under realistic conditions in section 5. Finally, section 6 summarizes the policy trie representation and discusses some areas of future research.

2 Firewall Security Policies

As previously described, a firewall security policy has been traditionally defined as an ordered list of firewall rules [21], as seen in table 1. A rule r can be viewed as an ordered tuple of sets [20], for example rule $r = (r[1], r[2], \dots, r[k])$. Each tuple $r[l]$ can be fully specified or contain wildcards ‘*’ in standard prefix format. For example the prefix $192.*$ would represent any IP address that has 192 as the first dotted-decimal number. For TCP/IP networks, rules are represented as a 5-tuple as seen in table 1. The tuples for TCP/IP are: protocol, IP source address, source port number, IP destination address, and destination port number. Order is necessary among the tuples since comparing rules and packets requires the comparison of corresponding tuples. In addition to the prefixes, each firewall rule has an action, which is to accept or deny. An accept action passes the packet into or from the secure network, while deny causes the packet to be discarded. Using the rule definition, a standard **security policy** can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, \dots, r_n\}$.

Similar to a firewall rule, a packet (IP datagram) d can be viewed as an ordered k -tuple $d = (d[1], d[2], \dots, d[k])$; however, wildcards are not possible for any packet tuple. An arriving packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. This is referred to as a first-match policy and is utilized by the majority of firewall systems [17]. A match is found between a packet and rule when every tuple of the packet is a proper subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \quad \text{iff} \quad d[l] \subseteq r_i[l], \quad l = 1, \dots, k$$

For example the packet $d = (\text{TCP}, 140.1.1.1, 90, 130.1.1.1, 20)$ would match the first and the fifth rules in table 1. However using the first match policy, the packet would be accepted

No.	Proto.	Source		Destination		Action
		IP	Port	IP	Port	
1	TCP	140.*	*	130.*	20	accept
2	TCP	140.*	*	*	80	accept
3	TCP	150.*	*	120.*	90	accept
4	UDP	150.*	*	*	3030	accept
5	*	*	*	*	*	deny
6	UDP	130.*	*	*	*	accept

Table 2: Example TCP/IP security policy where the sixth rule is shadowed by the fifth rule. Such anomalies are not considered in this paper.

since the first rule is the first match. Although the two rules match the same packet and have different actions, it is important to note that this is **not** considered an anomaly.

Shadowing is a type of anomaly, which occurs when a rule r_{i+k} matches a preceding rule r_i , thus rendering r_{i+k} obsolete. For example, table 2 gives a policy where the sixth rule is shadowed by the fifth ($r_6 \Rightarrow r_5$); thus, the new rule will never be utilized. Security policy anomaly detection and correction is the subject of continued research [2, 13, 21] and is not the focus of this paper. Therefore, this paper will assume such filter conflicts are not present in the security policies.

A security policy R is considered comprehensive if for every possible legal packet d a match is found using R . The policy given in table 1 is comprehensive due to the last rule. Furthermore, we will say two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. This definition will be extended to include different policy representations. As described in the introduction, this paper is interested in improving network firewall performance. Firewall performance will be measured using the number of *tuple-comparisons* required to find the first match. The worst case performance for a list-based representation is $k \cdot n$ tuple-compare, which occurs when the last rule (default rule) is the first match.

3 A New Security Policy Representation

In this section a new security policy representation called the *policy trie* is introduced, which provides faster processing of packets while maintaining the integrity of the original policy. The policy trie T is a n -ary trie structure consisting of k levels that stores a security policy. Each level $T[l]$ corresponds to a rule tuple (except for the root), while nodes on a certain level store the tuple values $T[l, v]$. Unlike the standard binary trie structure [1], the policy trie is unique since a node can have multiple children, similar to an n -ary tree. This is required since a node will store a rule tuple (multibit field), not just a single bit as done in [18]. Tuples at each level are organized from specific to general (reading left to right). For reference, levels will be numbered sequentially starting with zero for the root node. Likewise, nodes of a particular level will be numbered sequentially starting with zero for the left-most node. Since each level stores a tuple, a path from the root node to a leaf represents a firewall rule, as seen in figure 1.

To create a policy trie T , rules are added in the order they appear in R . A rule r is added to T by starting with the root node on the first level and comparing the values of its children with the

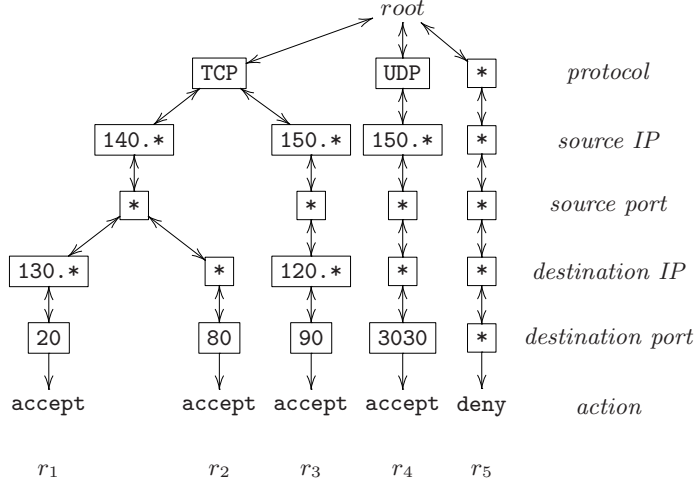


Figure 1: Policy trie representation of the firewall rules given in table 1.

corresponding tuple of r . If one of the children is equal (not just a subset) to the corresponding rule tuple, then r will share this node and a new node is not added for this level. The trie is traversed to that node and the process of comparing children to the rule is repeated using the next tuple in r . If an exact match does not exist, a new child node is created that contains the value of the corresponding packet tuple. In order to maintain the specific-to-general organization of the trie, the new node is inserted in the rightmost available position such that it is before (to the left of) any sibling that is a superset of the new node. The new node forms a chain of nodes that stores the remaining tuple values of the rule.

Consider the events that occur when rule r_2 is added to the trie given in figure 1. Rule r_2 has the same protocol, IP source, and source port values as r_1 ; therefore, r_2 will share these nodes. Since the destination IP is different, this forms a chain consisting of this tuple, the destination port, and action. This chain connects to the source port node, which adds r_2 to the trie. It is this structure that allows the elimination of multiple rules simultaneously. For example if a UDP packet is compared using the trie given in figure 1, then rules r_1 , r_2 , and r_3 are eliminated after the protocol tuple comparison.

Once the new rule is added, if any nodes exist to the right of the new rule, then this represents a rule re-order and may result in a shadowing. The intersection of the new rule and each of these right-most rules is taken and the action of right rule, the rule that appears first in the ordered policy, is applied.

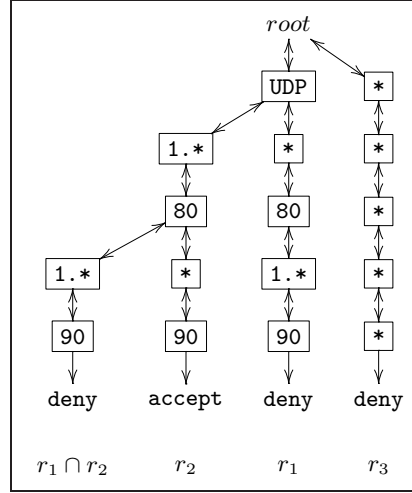
Definition The intersection of rule r_i and r_j , denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]), \quad l = 1, \dots, k$$

For all intersections that yield valid rules, the results form a subtree of the newly added rule. The same method is applied to this subtree. For example consider the rules given in figure 2(a). Note that the relative order of the rules must be preserved; otherwise integrity is not maintained. When r_2 is added to the policy trie, the source IP address will cause the rule to be placed before r_1 and

No.	Proto.	Source		Destination		Action
		IP	Port	IP	Port	
1	UDP	*	80	1.*	90	deny
2	UDP	1.*	80	*	90	accept
3	*	*	*	*	*	deny

(a) Example rule list, where the rules must maintain their relative order.



(b) Policy trie that requires intersection of r_1 and r_2 .

Figure 2: List and trie representation of a security policy where the policy trie requires the intersection operation to maintain the integrity of the list.

the intersection must be taken. The intersection of r_1 and r_2 is $(\text{UDP}, 1.*, 80, 1.*, 90)$. The result of the intersection indicates a packet can match both rules, for example $d = (\text{UDP}, 1.1.1.1, 80, 1.1.1.1, 90)$. Therefore this intersection rule, with the action of r_1 , must be added to the trie. The final policy trie is given in figure 2(b), which maintains the integrity of the policy given in 2(a). When r_3 is added, it is located to the right of r_1 and r_2 , thus intersections are not performed. In this example rules r_1 and r_2 are considered a *partial-match* [2]. The DAG policy representation described in [6] will not correctly represent partial-match rules, because subsets (as done with the match operation) are used to create the structure instead of intersections. As a result, the DAG structure described in [6] is **not** suitable for firewall policies since integrity cannot be maintained.

To process a packet d using the policy trie T (also referred to as searching T), the corresponding tuple of the packet is compared with the children of the root node. Comparisons of nodes are always performed from left and right, or specific to general. Once a match is found, the current node is marked and the trie is traversed to the matching child. The procedure is repeated with the remaining rule tuples. If no match is found, the search backtracks to the parent node and finds the next matching node that has not been visited, continuing the process of left to right comparison. Once a path has been found from the root node to a leaf where all the rule tuples match $(p[l] \subseteq T[l, i], l = 1, \dots, k)$ the associated action is performed.

3.1 Policy Trie Integrity

As previously stated, a necessary objective of any policy representation is its ability to maintain the policy integrity. This occurs if the new representation is equivalent to the original list-based

policy. A policy trie T is equivalent to the original security policy R for any legal packet d , if searches of T and R result in the same action being performed. Unlike other representations, this is proven true in the following theorem.

Theorem 3.1 *A policy trie T is equivalent to the original security policy R .*

Proof Assume a trie T is constructed with k levels using the process previously described from an n rule security policy (ordered list). Furthermore assume the completed trie is searched using the previously described method. One of the following three cases will occur during the creation of the trie.

Case 1 *Rule reorder does not occur during creation.* If rule reorders do not occur during the creation of T , then rules will appear from left to right in T as they appear in R (an in-order traversal of T yields R). As a result, the policy representations are equivalent because nodes are tested from left to right.

Case 2 *Rule reorders occur without intersections.* Consider a trie T consisting of $n - 1$ rules from R , which are added using the process previously described. In addition, let the $n - 1$ rules be ordered in T as they are in R . Assume a new rule r_n is added to T and is located to the left of an existing rule r_m . Let S be the set of rules that appear to the right of r_n in T , $S = \{r_i, m \leq i < n\}$. If the rules in S do not intersect with r_n , then a packet cannot match both r_n and any rule in S . As a result, testing r_n before any rule in S is not significant since shadowing is not introduced; thus, the reorder does not effect policy integrity.

Case 3 *Reorder and intersections occur during creation.* Consider a trie T consisting of $n - 1$ rules from R , which are added using the process previously described. In addition, let the $n - 1$ rules be ordered in T as they are in R . Assume a new rule r_n is added to T and is located to the left of the existing rule r_m . Let S be the set of rules that appear to right of r_n in T , $S = \{r_i, m \leq i < n\}$. Assume r_n does intersect with rule r_i in S . The intersection represents the set of packets that match r_n and r_i , where the tuples of the intersection are the more specific of the two rules. There must be at least one tuple in r_i more specific than the corresponding tuple in r_n , otherwise r_n is shadowed in the original rule list. The intersection rule will be located to the left of r_n and have the action of the r_i . Therefore, the intersection rule will always be tested first, and if true the action of the rule r_i is applied. This is the correct response; therefore, the reorder will not affect integrity.

■

3.2 Push-Down Policy Tries

The policy trie as described thus far may require *backtracking* when a packet is processed (search is performed). Backtracking searches can have a worst-case performance that is equal to a list representation [16]. Although the penalty for backtracking in an n -ary trie is not as severe as a standard binary version, the conversion to a non-backtracking trie can reduce the number of tuple-compare, which is the objective of the representation.

A non-backtracking policy trie, referred to as a push-down policy trie, is created by replicating, or *pushing-down* general rules in the original policy. A general rule is a superset of at least one other rule in the policy, and is defined as a range of values, containing at least one wild-card in the standard prefix notation. The push-down procedure replicates more general rules in subset

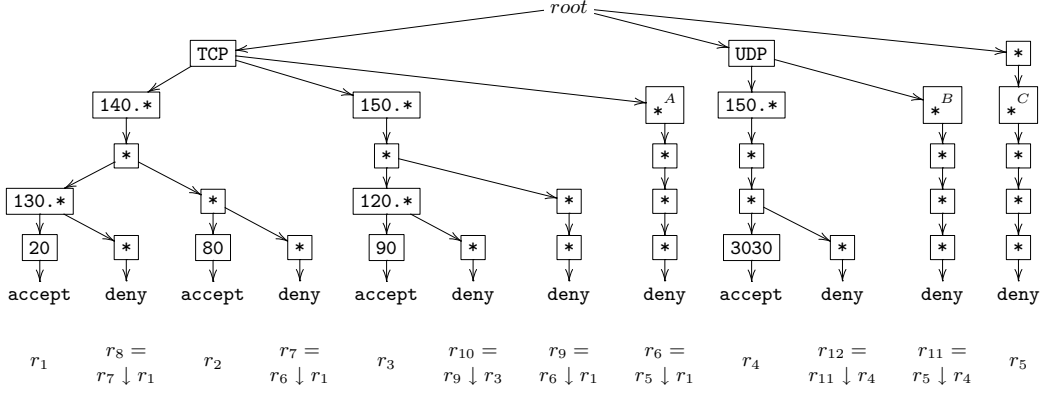


Figure 3: Push-down policy trie representation of the firewall rules given in table 1.

subtries, that would match the same packets as the general rule. As a result, the union of the push-down rules is a proper subset of the original rule.

Definition The push-down of rule r_g to r_s , denoted as $r_g \downarrow r_s$ is

$$r_g \downarrow r_s = (r_s[1, \dots, l], r_g[(l + 1), \dots, k], r_g[action])$$

where the $r_g[i] = r_s[i], i = 1, \dots, (l - 1)$ and l is the index of the first tuple of r_s that is a subset of the corresponding tuple in r_g .

A general rule can only be pushed-down to rules that appear to left of it in the trie. Furthermore, a non-backtracking policy trie is created when **all** general rules are pushed-down; therefore, $r_g \downarrow r_s, \forall s < g$. This can be easily implemented using a post-order traversal of the policy trie. Note that while push-down always creates a rule that is more specific than the general rule being pushed, the resulting rule may still be a superset of other rules in the policy trie, and thus must be pushed down. For example, consider the push-down policy trie given in figure 3. When rule r_5 is pushed-down to rule r_1 it creates rule r_6 . This is a general rule that is pushed-down again to rule r_1 yielding r_7 . This process repeats yielding r_8 , which cannot be pushed-down further.

As done with the original (backtracking) policy trie, we must be certain that the integrity of the original policy is maintained when using a push-down policy trie. Theorem 3.2 proves that push-down and original policy tries are equivalent. Therefore it can be stated that the push-down policy trie maintains integrity since, the original policy trie can be proven to do so.

Theorem 3.2 A push-down policy trie T_p is equivalent to the original policy trie T , which is equivalent to the original security policy R .

Proof Assume a push-down trie T_p is constructed with k levels from an n rule ordered list. Consider node i on level l , $T_p[l, i]$, that is part of rule r_s . The children of node i include the children of siblings (of node i) that appear to the right, if node i matches the sibling. Recall the push-down procedure is recursive. If a tuple value is not present among the children of node i , the associated rule(s) are not a possible alternative since node i is not a match. Furthermore, given the procedure for constructing the push-down trie, the children are always ordered as they appear on the right. The rules will be tested in T_p in the same order as T which is equivalent to R . ■

3.3 Worst Case Analysis

As described in the previous section, the push-down policy trie offers a performance increase by eliminating the need to traverse backwards. In this section, the worst case performance and storage requirement of the push-down trie is analyzed theoretically. A primary objective of the policy trie representation is to reduce the number of tuple-comparisons required per packet. Before this can be done, two important lemmas about push-down policy tries are required.

Lemma 3.3 *A node in a push-down policy trie cannot have more than n (the number of rules) children.*

Proof Every child of a node is unique, duplicates are not possible. If a node has n children then all possible set values for the rule set are present for the tuple (each level represents a tuple). The intersection and push-down operations never introduce a new set value for a tuple as the operations create a new combination of tuple values; therefore, the result of the either operation would share one of the existing children if the node already has n children. ■

Lemma 3.4 *Each node traversal at a particular level in a push-down policy trie eliminates at least one rule from consideration.*

Proof Every child of a node must be unique, and every child value is an actual tuple value from the rule set. Consider node i in the push-down trie. A rule that equals the value of node i cannot also be a part of any other rightward sibling of node i without either duplicating the node or violating the specific to general creation of the trie. This is evident since push-down and intersection operations never result in a rule being replicated towards the right. Furthermore, new push-down and intersection nodes are always located to the left of the leftmost node involved in the operation. Combining the left-to-right property of the search with the non-backtracking nature of the push down trie, it can be said the rule in question is eliminated from consideration. ■

The previous two lemmas provide important bounds on the structure of any push-down policy trie. As a result, the worst case number of tuple-comparisons is $O(n+k)$, which is proven in theorem 3.5. Comparing this bound with the worst case for a list-based representation, the push-down policy trie requires a fraction $(1/k)$ of the processing.

Theorem 3.5 *A comprehensive push-down trie consisting of k levels and constructed from n rules requires $O(n+k)$ number of tuple-comparisons to match a packet in the worst case.*

Proof We must always traverse every level of the trie (k tuple-compares) to determine the action. Following from lemmas 3.3 and 3.4, we know that traversing a node eliminates at least one rule, which would require n traversals in the worst case. As a result, the worst case number of tuple-compares is $k+n$. ■




Figure 4: Worst case push-down policy trie ($r_1 \Rightarrow r_2$, $r_1 \Rightarrow r_3$, and $r_2 \Rightarrow r_3$) and the DAG equivalent. Assume the original security policy has three rules, where a rule has only three tuples.

4 Storage

This paper has focused on reducing the number of comparisons needed to determine the appropriate match for a given packet and security policy; however, the amount of storage required for the different policy representations is also important. Although the amount of storage required for a trie would appear at first to be lower than a list, intersection and push-down operations do increase the number of nodes in the push-down trie. This is evident in the number of nodes required for the push-down trie depicted in figure 3 as compared to the original trie given in figure 1.

Given an n rule firewall policy, push-down causes the worst case storage requirement to occur under two specific conditions. The first condition is $r_i \Rightarrow r_j, \forall i < j < n$, a rule matches all the rules that appear to the right, maximizing the number of push-downs that occur. The second condition is $r_i[l] \neq r_j[l], \forall i < j < n, 1 \leq l \leq k$; none of the tuples are equal and nodes are never shared. This worst case is depicted in figure 4(a), where a 3 rule list requires 20 nodes in the push-down policy trie. However, the number of nodes required by the trie can be greatly reduced by converting it into a Directed Acyclical Graph (DAG) [1, 16]. In the context of a DAG, the push-down operation directly references nodes instead of replicating them as in the policy trie. For example, consider the push-down trie given in figure 3. The parents of the nodes labeled A and B could point to the node labeled C , which eliminates the need for new nodes for rules r_6 and r_{11} . The other push-down rules can be replaced in a similar fashion. Similarly, the DAG equivalent of the push-down policy trie given in figure 4(a) is depicted in figure 4(b) and requires significantly fewer nodes. The DAG conversion causes the worst case push-down trie to only require $k \cdot n$ tuples, which equals the storage requirement for a list representation.

The intersection of two rules, required when rule reordering occurs, may also result in a *new rule* (a new combination of existing tuples). The worst case policy would require the intersection among all rules to result in a valid rule, where tuples of the new rule alternate between the two rules. In addition, the rules would have to be listed according to the first tuple from most general to most specific. In this situation, intersection operations result in chains that cannot be replaced by DAGs and the worst case number of nodes required to store the policy trie would be $O(n^2)$. Although this is a higher space requirement than the standard list representation, it only occurs under very

No.	Source IP	Destination IP	Action
1	*	20.20.*	accept
2	10.*	20.*	drop
3	10.10.*	*	accept

Table 3: An abbreviated TCP/IP security policy that causes the worst case number of intersections.

specific circumstances. Furthermore, the significance of the additional space requirement is relative to the frequency of the worst case packet(s).

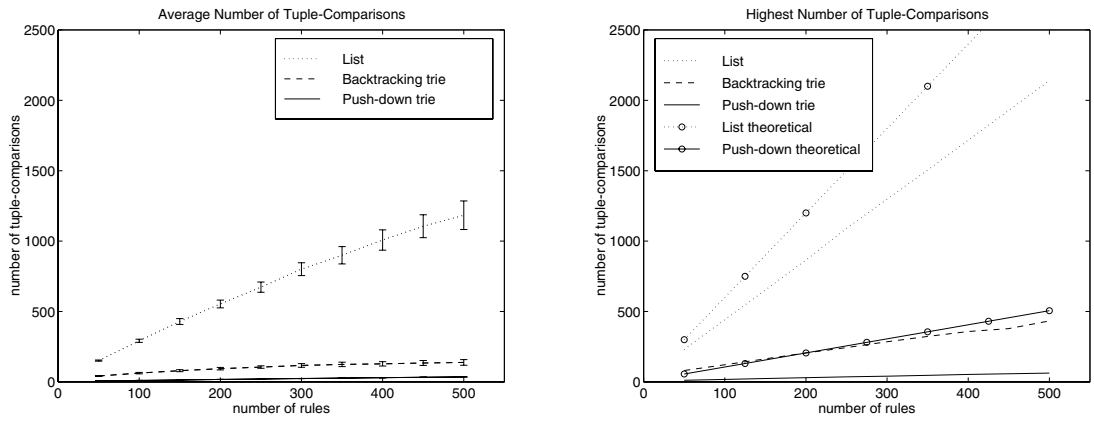
Theorem 4.1 *A policy trie will require $O(n^2)$ tuple storage in the worst case.*

Proof Assume the values for tuples in policy R are proper subsets of each other. The maximum number of intersections occur when the rules of policy R are ordered such that the values for the first tuple appear from general to specific. Assuming the first level of the policy trie T stores the first tuple in R , then T will store the rules in reverse order and the intersection of rule i must be taken with rules $i + 1$ through n . Note, the values of at least one tuple must occur from specific to general in R , otherwise shadowing occurs. Therefore, assume the values for the second tuple in R are arranged from specific to general, while the values for the third tuples are arranged from general to specific. Assume this alternating order occurs for the remaining tuples, as seen in table 3. As a result, the intersection operation creates a chain of tuples, where the chains cannot be replaced with a DAG since it has a unique order of tuple values. The number of tuples required for T is $k \cdot n$ for the original rules and $(k - 1) \sum_{i=0}^{n-1} i$ for the intersections. Therefore, the total number of tuples required in the worst case is $k \cdot n + (k - 1) \cdot n \cdot (n - 1)/2 \approx O(n^2)$ ■

5 Experimental Results

The previous section described a new network security policy representation called a policy trie that was shown to provide theoretically better performance than the standard list-based representation. Simulation results presented in this section will confirm the worst case number of tuple-comparisons and show that similar performance gains are achieved in the average case. In addition, the average and worst case storage requirements for the different policy representations are presented and will be shown to also remain within their theoretical bounds.

Simulations were conducted using list, backtracking trie (original trie), and push-down trie representations of firewall policies. A random rule generator was used to create valid rule sets with a realistic degree of rule intersection. The generator was set to allow a slightly lower number of tuple permutations at high levels (source, source port, etc.) so that the shape of resulting policy tries would mirror those of real-world firewall rule sets. Policy sizes ranged from 50 rules to 500 rules, where 50 different policies were generated per policy size. Sets of 10,000 packets were passed through representations of each policy and the resulting decisions made were validated against the original rule set. Statistics concerning the average and worst case number of tuple-comparisons were recorded as well as the amount of storage required for each policy representation



(a) Average number of tuple-comparisons. Error bars represent one standard deviation. (b) Highest number of tuple-comparisons.


Figure 5: The average and worst case number of tuple-comparisons required for the firewall experiments. Push-down trie provided the best performance in both cases.

5.1 Tuple-Comparisons Results

Results for the tuple-comparisons are given in figure 5. As expected, when the policy sizes increased, both trie representations always performed considerably better than the linear rule set. In all cases, each representation reached the same decision, indicating that they were equivalent; thus maintaining integrity. As seen in figure 5(a), the average performance for backtracking tries appears to be similar to that of push-down tries. However, the backtracking trie required 5 times as many comparisons on average than the push-down trie, while the original list required 34 times as many tuple-comparisons on average.

The variance for the average number of comparisons in push-down tries was slightly lower than that of backtracking tries. As a result, push-down tries sometimes performed significantly better than their backtracking counterparts. Compared to linear implementations, the variance of trie-based implementations was very low and relatively constant. The standard deviation of the average number of comparisons in either trie implementation never rose above 20 per packet over 10,000 packets. Though this number may seem significant, the variance of linear policies averaged more than 46 comparisons and sometimes ranged as high as 100 comparisons.

The worst case number of tuple-comparisons required by each representation is given in figure 5(b). Similar to the average case results, both trie representations out-performed the list representation. Compared to the push-down trie the backtracking trie required 7 times as many of tuple-comparisons to reach a decision in the worst case, while the list representation required 31 times as many. In addition, the performance of push-down tries and list representations were within the theoretical bounds.



(a) Average number of tuples stored. Error bars represent one standard deviation.

(b) Highest number of tuples stored.

Figure 6: The average and worst case number of tuple stored for the firewall experiments. Backtracking trie required the least amount of storage in both cases.

5.2 Storage Results

The amount of storage required, measured in the number of tuples, is depicted in figure 6. As predicted, when policy sizes increased the backtracking tries consistently used less (a third for these experiments) of the storage required by the list representation. In contrast, the push-down tries required the most storage. The push-down trie storage was nonlinear with respect to the number of rules and on average required 10 times as much storage as the backtracking trie. Furthermore, the variance of push-down tries averaged over 1,000 nodes, while the variance of backtracking tries averaged less than 50 nodes. Although the push-down trie representation required the most storage, the observed worst case requirement was well below the theoretical upper bound of n^2 , requiring on average 92% fewer tuples.

The amount of storage required by both trie representations is directly related to the degree of rule overlap. Backtracking tries benefit from rule overlap because overlap results in a greater number of shared nodes. In contrast, the storage requirement for push-down tries improves when there are fewer subset relations in a policy, since fewer push-down operations occur.

6 Conclusions

Network firewalls must continue to match, or exceed, the ever increasing speed and volume of network traffic if they are to remain effective. Unfortunately, the traditional single machine firewall that utilizes a list-based security policy can easily become overwhelmed in this environment. Therefore, new methods to increase firewall performance are needed to meet the demands of the next generation of networks and applications.

This paper introduced a new firewall security policy representation called the policy trie, that

maintains policy integrity while requiring significantly less processing time. Rules are modeled as an ordered set of tuples, allowing the precedence relationship between rules to be maintained. The security policy is represented as an n -ary retrieval tree (trie), which combines the fast search properties of a trie with the flexibility of a n -ary tree. The policy trie stores similar rules together, which allows the elimination of multiple rules as a packet is processed. This is in contrast to a standard list-based representation, which can only consider rules individually. This yields a worst case performance for the policy trie that is only $1/k$ of a list representation, where k is the number of tuples in a rule. Furthermore unlike other representations, the policy trie maintains the integrity of the original policy. The integrity and performance improvement of the policy trie was proven theoretically and demonstrated using simulation under realistic conditions.

While the policy trie has shown great promise, more research is needed to evaluate its ability to manage a dynamic policy. In many cases, rules must be added and removed over time to reflect the current security status (for example, stateful firewalls used for connection tracking [21]). A simple solution would just recreate the trie from the corresponding list-based policy once a rule is removed. Another approach would tag rules if they are the result of intersection or push-down operations, which allows their quick removal if the corresponding original rule is removed. Another area of research would investigate the average performance of the firewall system. Possible methods for improving the average case would involve sorting the trie nodes based on match probabilities or reorganizing the tuple order of the trie. This must be done while maintaining the integrity of the policy, which is not a trivial problem.

Acknowledgement

The authors of this paper would like to thank Patrick Wheeler at The Department of Computer Science at U.C. Davis for his valuable review and input.

Appendix: Policy Trie Operations

Two important policy trie operations are adding a new rule and searching. Both operations are described in this section using a C++ pseudocode; therefore objects are used to represent a packet, a firewall rule, and the policy trie (which consists of trie nodes). The rule and packet objects referenced by the functions store tuples, while the data structure used for the policy trie object is an n -ary retrieval tree (trie). Furthermore, the functions assume operator overloading has been correctly implemented for the objects, which includes intersection.

As described in section 3, the process of adding a new rule to the policy trie requires maintaining the integrity of the original rule list. As previously described, tuples are always stored from specific to general. If this causes a rule to be placed out of order, with respect to the original list, intersections must be taken to maintain integrity. This process is given on the next page.

```

1 void addRule(PolicyTrieNode* nPtr, Rule newRule, int tuple)
2 begin
3     match = false
4     for i = 0; i < nPtr->numChildren AND !match; i++
5         if newRule[tuple] == nPtr->child[i]->value then
6             match = true
7             if (tuple + 1) < maxTuples then
8                 addRule(nPtr->child[i], newRule, tuple + 1)
9             endif
10        endif
11    endfor
12
13    // add newRule (the remaining chain)
14    if !match AND tuple < maxTuples then
15        nPtr->child[nPtr->numChildren] = new PolicyTrieNode(newRule, tuple)
16        nPtr->numChildren++;
17    else
18        newRule already present ...
19    endif
20
21    // intersection with the right-most siblings
22    for ; i < nPtr->numChildren; i++
23        if match then
24            // intersect with rules containing this tuple
25            if c->child[i]->value  $\cap$  newRule[tuple]  $\neq \emptyset$  then
26                intersect (nPtr->child[i], newRule, tuple + 1)
27            endif
28        endif
29    endfor
30 end
31
32 void intersect (PolicyTrieNode* nPtr, Rule newRule, int tuple)
33 begin
34     for i = 0; i < nPtr->numChildren; i++
35         // intersect with rules containing this tuple
36         if c->child[i]->value  $\cap$  newRule[tuple]  $\neq \emptyset$  then
37             intersect (c->child[i], newRule, tuple + 1)
38         endif
39     endfor
40 end

```

As described in section 3, searching the policy trie for the correct match can be done recursively. The search is a *left-to-right* process, since tuples are stored from specific to general order. The pseudocode for searching the policy trie is given on the next page.

```

1  bool searchTrie(PolicyTrieNode* nPtr, Packet pkt, int tuple, RuleAction& action)
2  begin
3      if tuple > maxTuples then
4          action = nPtr→action
5          return true
6      endif
7      match = false
8      for i = 0; i < nPtr→numChildren AND !match; i++
9          if pkt[tuple] ⊆ nPtr→child[i]→value then
10             match = searchTrie(nPtr→child[i], pkt, tuple + 1)
11         endif
12     endfor
13     return match;
14 end

```

References

- [1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *Data Structures and Algorithms*. Addison-Wesley, 1987.
- [2] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.
- [3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.
- [4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.
- [5] M. Christiansen and E. Fleury. Using Interval Decision Diagrams for Packet Filtering. Technical report, BRICS, 2002.
- [6] D. Decasper, Z. Dittia, G. Parulkar, and B. Plattner. Router Plugins: A Software Architecture for Next-Generation Routers. *IEEE/ACM Transactions on Networking*, 8(1), February 2000.
- [7] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. Small Forwarding Tables for Fast Routing Lookups. In *Proceedings of ACM SIGCOMM*, pages 4 – 13, 1997.
- [8] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In *Proceedings of INFOSEC'COM*, 1998.
- [9] A. Feldmann and S. Muthukrishnan. Tradeoffs for Packet Classification. In *Proceedings of the IEEE INFOCOM*, pages 397 – 413, 2000.
- [10] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

- [11] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.
- [12] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.
- [13] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.
- [14] HiPAC. High Performance Packet Classification. <http://www.hipac.org>.
- [15] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In *Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000*, 2000.
- [16] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.
- [17] V. P. Ranganath and D. Andresen. A Set-Based Approach to Packet Classification. In *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 889–894, 2003.
- [18] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and Scalable Layer Four Switching. In *Proceedings of ACM SIGCOMM*, pages 191–202, 1998.
- [19] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.
- [20] P. W. Zehna. *Sets with Applications*. Allyn and Bacon, 1966.
- [21] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.